
Altitude Deviations: Breakdowns of an Error Tolerant System

Everett A. Palmer, Ames Research Center, Moffett Field, California
Edwin L. Hutchins, University of California, San Diego, La Jolla, California
Richard D. Ritter, Ames Research Center, Moffett Field, California
Inge vanCleemput, Wellesley College, Wellesley, Massachusetts

December 1991

National Aeronautics and
Space Administration

U.S. Department of Transportation
Federal Aviation Administration

Table of Contents

Summary.....	1
Introduction.....	1
Aviation Safety Reporting System.....	3
Approach.....	4
Summary of Results.....	5
Sample Narratives from ASRS Reports	8
1. Crew Distracted by Programming the Flight Management System.....	8
2. Error in Programming Flight-Management System.....	10
3. A Missing Minus Sign.....	11
4. The Green Arc.....	11
5. An Unexpected Clearance and a Tired Crew.....	12
6. One-Pilot-off-Radio Incident.....	13
7. Lack of an Aural Approaching-Altitude Alert Blamed.....	13
8. High Performance Low Altitude Level Off.....	14
9. Ambiguous Shorthand and a Misset Altimeter.....	15
10. Misplaced Trust in the Autoflight System.....	16
Field Study Findings.....	17
Building a System Description.....	17
Terms of Task Analysis.....	18
System of Representations.....	18
Theory of Action.....	18
Chronology of Altitude-Change Task Elements.....	20
Hear and Record the Clearance.....	21
Understand the Clearance and Decide on Acceptance.....	22
Read Back the Clearance.....	23
Create Machine-Readable Description of Flight Path.....	24
More Complex Machine-Readable Descriptions of Flight-Path.....	26
Select Mode of Operation.....	29
Couple the Autoflight Computers to the Flight Controls.....	29
Supervise the Autoflight System.....	30
Discussion.....	31
Feedback and Error Tolerance.....	32
Distributed Access to Information.....	32
Distractions.....	33
Intelligent Feedback.....	34
Recommendations.....	34
Procedures.....	34
Training.....	34
New Equipment.....	
Datalink.....	35
Future Design Possibilities.....	35
Continuing Research	35
Conclusion.....	36
Appendix: Glossary.....	39
References.....	41

Altitude Deviations: Breakdowns of an Error-Tolerant System

Everett A. Palmer, Edwin L. Hutchins,^{*}
Richard D. Ritter, and Inge vanCleeemput[†]

SUMMARY

Pilot reports of aviation incidents to the Aviation Safety Reporting System (ASRS) provide a window on the problems occurring in today's airline cockpits. The narratives of 10 pilot reports of errors made in the automation-assisted altitude-change task are used to illustrate some of the issues associated with pilot and automatic-systems interactions. These narratives are then used to construct a description of the cockpit as an information processing system. Describing the cockpit as a single information processing system is useful because the system behaviors of interest are not determined solely by the behavior of the humans in the system. Altitude deviations in transport aircraft are usually the result of several small problems. The information processing analysis also highlights the variety of languages and media used in the cockpit to describe the flight path as clearance information is processed by the cockpit system. The analysis concentrates on the error-tolerant properties of the system and on how breakdowns can occasionally occur. An error-tolerant system can detect and correct its internal processing errors. The cockpit system consists of two or three pilots supported by autoflight, flight-management, and alerting systems. These humans and machines have distributed access to clearance information and perform redundant processing of information. Errors can be detected as deviations from either expected behavior or as deviations from expected information. Breakdowns in this system can occur when the checking and cross-checking tasks that give the system its error-tolerant properties are not performed because of distractions or other task demands. The report concludes with recommendations based on the analysis for improving the error tolerance of the cockpit system.

INTRODUCTION

One objective of NASA's Aviation Safety/Automation program is to develop error-tolerant systems that can detect errors and correct them before they result in undesired consequences. In attempting to design systems that are more error-tolerant, it is useful to examine and account for the impressive error-tolerant properties of today's aviation system. This report examines the error-tolerant properties of the very common aviation task of changing altitude. Our approach is to study pilot reports to the Aviation Safety Reporting System (ASRS) of breakdowns in this normally error-tolerant task. These reports are used to assemble a composite description of the flight deck as an information processing system. We use the system description to show why the current system is so robust in the face of seemingly high error rates. The report emphasis is on the effect of increased cockpit automation—for aircraft navigation, guidance, and control in the current generation of glass cockpit airliners—on the type of altitude deviations reported by pilots.

Aircraft in controlled airspace are cleared by air traffic controllers to fly at specific altitudes. Pilots are given clearances such as "NASA 1003, descend and maintain one one thousand feet." On accepting this clearance, the pilot is expected to descend without delay to 11,000 ft and then maintain 11,000 ft. It is important for pilots to adhere accurately to altitude clearances because a

^{*} Professor, University of California, LaJolla, California.

[†] Ames Associate, Wellesley College, Wellesley, Massachusetts.

primary method air-traffic controllers use to keep aircraft from colliding is to assign them to different altitudes. These assigned altitudes are usually at least 1000 ft apart.

The three types of altitude deviations are (1) an *undershoot*, in which the aircraft levels off before reaching the assigned altitude; (2) an *overshoot*, in which the aircraft levels off after reaching the assigned altitude; and (3) a *deviation* in which the aircraft levels off at the assigned altitude but later deviates from the assigned altitude without clearance. The tolerance for achieving and maintaining an assigned altitude is 300 ft. Another kind of altitude clearance is a crossing restriction. For example, “NASA1003, cross SUNOL at or below eight thousand feet.” An altitude deviation can also occur if an aircraft does not climb or descend quickly enough to meet a crossing restriction. Pilots can be cited by the FAA for failing to adhere to a clearance and can receive a warning, have a fine imposed, or have their licenses suspended.



Figure 1. Typical glass-cockpit displays and controls used in the altitude-change task.

A surprisingly complex system of people and equipment is involved in performing and monitoring the task of changing altitude¹ in an airline transport aircraft. The altitude-change task usually begins with a radio call from the air-traffic controller to the aircraft crew requesting that they change to a new altitude. One pilot—usually the pilot-not-flying (PNF)—handles the task of

¹ A blow-by-blow description of pilot activity during the initial part of a normal altitude change is contained in Hutchins and Klausen (1992).

communicating with air-traffic control (ATC). The PNF receives the new clearance and reads it back to the air-traffic controller who monitors the “read-back” for communication errors; one of the pilots then enters the new altitude into the altitude alerter.² The function of the altitude alerter is to alert the flight crew when the aircraft is approaching the target altitude and to alert the crew of deviations from the target altitude of more than a few hundred feet. All alerts are signaled by a light next to the altimeter. Deviations from the specified altitude also trigger an aural alert. Some aircraft also have an aural alert that sounds when the aircraft climbs or descends to within 900 ft of the specified altitude in order to remind the crew that the aircraft is approaching the new altitude.³

The value set in the altitude window is also the target altitude for the autopilot. If the new altitude is lower than the aircraft’s present altitude, the pilot flying can manually reduce engine power and initiate the descent, or use the autoflight system⁴ to make the altitude change. In either case, the pilots should remain aware of altitude during the climb or descent to the new altitude. At 1,000 ft before the new altitude, a typical airline procedure is for the PNF to make a verbal call of “one thousand to go” or “out of twelve for eleven.” At 900 ft before the new altitude, the altitude-alert light illuminates. Approaching the new altitude, the pilot-flying (PF) or the autoflight system levels the aircraft and adjusts engine power to maintain speed. The controller can monitor the aircraft’s altitude on the ground radar display. The air-traffic controller also has a system that monitors each aircraft’s altitude and alerts the controller to deviations from assigned altitudes⁵.

Figure 1 illustrates the typical controls and displays in a glass-cockpit airliner that can be used during an altitude change; Figure 2 is a time-line of typical events. Figure 3 is a diagram of the possible paths over which information can flow between the people and equipment that make up the system that performs the altitude-change task.⁶ Breakdowns and errors in changing altitude will be illustrated with pilot narratives from the ASRS.

AVIATION SAFETY REPORTING SYSTEM

The Aviation Safety Reporting System (ASRS) is an incident reporting system run by NASA for the FAA. ASRS reports are submitted voluntarily by pilots and controllers. Pilots have two principal motivations for submitting ASRS reports describing errors and incidents in which they have been involved: first, responsible pilots may have a desire to report a safety problem so that it can be avoided in the future; and second, the FAA has agreed to grant pilots who file reports limited immunity from legal actions that might have been brought as a consequence of the incident described in the report. Since all reports are voluntarily submitted, there are fundamental problems with any attempt to conduct statistical analyses of the reported incidents. About all that can be concluded with certainty is that these reports indicate that at least some minimum number of different kinds of incidents have occurred. The key strength of ASRS incident reports is that they provide examples of the kinds of errors and incidents that are occurring in the aviation system. The ASRS office states in its standard cover letter that accompanies all reports that “We believe that the real power of ASRS lies in the report narratives. Here pilots, controllers, and others tell us about aviation safety incidents and situations in detail. They explain what happened, and more importantly, why it happened.”

For this study, a request was made to the ASRS office for reports⁷ of altitude-deviation incidents from airline pilots of conventional-technology and the newer, more automated glass-cockpit aircraft. Examples of the type of aircraft in the conventional cockpit sample are the B-727, B-737-100, B-737-200, B-747-200, DC-9, and DC-10. Examples of the type of aircraft in the

² See Federal Aviation Regulation (FAR) 91-165 for specifications of the altitude-alerter.

³ Altitude alerters in Boeing airplanes such as the B-757/767 signal the crew 900 ft before the target altitude; in Douglas airplanes such as the DC-9/MD-80 the alert margin is 700 ft; in Airbus aircraft the margin is 700 ft.

⁴ The term “autoflight” refers to the combined autopilot and autothrottle system.

⁵ This ground system is part of the “Operational Error Detection Program.” Pilots commonly refer to this system as The Snitch.

⁶ See also Degani, Chappell, and Hayes (1991) for a similar diagram of information flow.

⁷ ASRS incident reports are publically available.

glass cockpit category are the B-747-400, B-757, B-767, MD-88, A-310, and A-320. These glass-cockpit aircraft are equipped with flight-management computer systems and sophisticated autoflight systems. The ASRS incident reports do not provide enough information to identify reliably the exact make and model of the aircraft. In this paper, we will identify them only as conventional or glass-cockpit aircraft.




Figure 2. Time-line of events during the altitude-change task.

APPROACH

A one-page coding form was developed that allowed the ASRS incidents to be described in terms of the factors that were present in the incident. In the initial part of this study, this form was completed for altitude-deviations reports from the pilots of 50 conventional-cockpit and 50 glass-cockpit aircraft. The factors and their definitions are shown in Table 1.

Examples of ASRS narratives that illustrate the trends observed in this collection of altitude-deviations incidents are provided in a subsequent section. They are described in terms of recent results in the fields of human-computer interaction and cognitive science. The incidents are then used to construct a descriptive model of the information flow in the error-tolerant human-machine system that performs and monitors the altitude-change task. The report concludes with recommendations for procedures for the altitude-change task and for future equipment design.

SUMMARY OF RESULTS

Like the results of other ASRS studies of altitude deviations (Factors... 1977; Billings and Cheaney 1981), the factors in this sample of altitude-deviation incidents could be attributed either to a failure in information transfer or to a distraction from checking and cross-checking tasks. There were also several altitude-deviation reports that mentioned or implied that inexperience in operating the autopilot or flight-management systems was a factor. Table 2 presents a count of the factors that contributed to each reported altitude-deviation incident in the sample. Up to five of the factors in Table 1 were assigned to each incident.

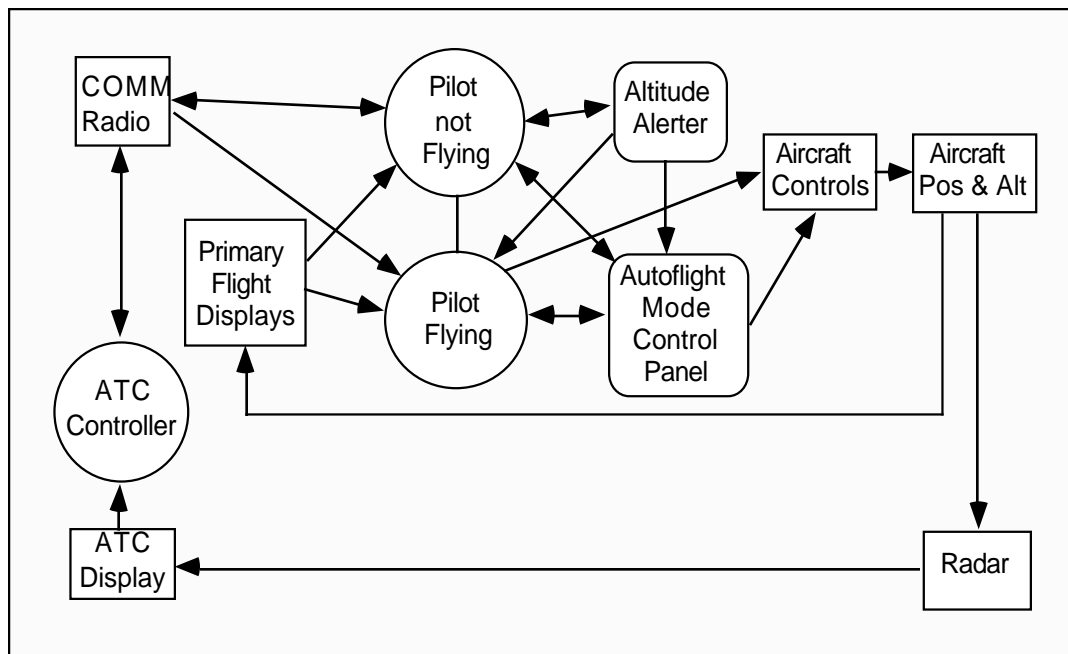


Figure 3. Potential information flow paths during the altitude-change task

Table 1. Factors Used to Describe Altitude-Deviation Incidents

CDU/FMS Programming.....	a problem programming the FMC
Emergency System.....	aircraft system malfunctioning
In-Flight Training.....	one pilot training another pilot
Lack of Knowledge.....	problem caused by lack of knowledge about system operation
High Workload.....	distraction caused by high workload
Complex Flight Path.....	problem owing to the complexity of the flight path
Aircraft System Problem.....	malfunction of some aircraft equipment
Weather.....	weather factor
Complacency.....	mentioned being complacent
Not monitoring autopilot.....	mentioned not monitoring autopilot
No cross-check.....	mentioned cross-checks not performed
Overfamiliar crew.....	mentioned being lax because of crew familiarity
Call-sign confusion.....	communication confusion owing to similar call signs
Read-back/hear-back.....	incorrect read-back to ATC not detected by ATC
Misunderstood.....	ATC clearance not understood
Pilot off radio.....	one pilot not monitoring ATC frequency
Frequency congestion.....	difficulty communicating with busy ATC
Crew coordination.....	crew not working well together
Altitude alerter misset.....	wrong altitude set in altitude-alerter window
Barometer misset.....	barometric pressure misset in altimeter
Different AC systems.....	confusion owing to different systems in different aircraft
Autopilot capture killed.....	automatic altitude capture disabled because of inadvertent action
Inexperience.....	lack of experience with autoflight/FMS
Complex autopilot.....	mentioned complexity of autopilot
Complex FMS.....	mentioned complexity of FMS
Loss of flight skill.....	mentioned loss of basic flight skills
Radio discipline.....	not using standard phraseology
Computer malfunction.....	autoflight system malfunctioned or did not work as expected
Long day.....	mentioned long duty time and/or fatigue as problem

Table 2. Descriptive Factors Assigned to Altitude-Deviation Reports

<i>Assigned factors</i>	<i>Aircraft cockpit type, number of reported occurrences</i>	
	<i>Traditional</i>	<i>Glass</i>
Distraction		
CDU/FMS programming	0	14
Emergency system	3	0
In-flight training	1	2
Lack of knowledge	3	5
High workload	7	10
Complex flight path	1	7
System problem	2	5
Weather	5	2
Complacency		
Not monitoring autopilot	6	10
No cross-check	8	13
Overfamiliar crew	2	1
ATC Comm		
Call-sign confusion	3	1
Read-back	21	12
Misunderstood	4	1
Pilot off radio	11	9
Frequency congestion	5	2
Cockpit Comm		
Crew coordination	19	14
Altitude alerter misset	0	5
Barometer misset	0	2
Training		
Different AC systems	1	3
Autopilot capture killed	1	4
Inexperience	3	15
Complex autopilot	0	5
Complex FMS	0	4
Loss of flight skill	1	3
Radio discipline	4	0
Fatigue		
Long day	6	10

SAMPLE NARRATIVES FROM ASRS REPORTS

The following ten ASRS narratives were chosen to illustrate the range of problems identified in the initial survey. They provide insight into how the new technology in the glass cockpit airliners has changed the nature and types of errors in the altitude change task. In the following section, these incidents will be used to construct a composite descriptive model of information flow and processing during the altitude change task.

1. Crew Distracted by Programming the Flight Management System

The following report⁸ from the first officer of a glass-cockpit airliner describes an altitude deviation that occurred during the descent phase of flight as both pilots attempted to program a holding pattern into the flight-management system (FMS).

Descending from higher flight levels to 15,000 ft on Center clearance had anticipated and received clearance to hold at BUCKS due to anticipated weather delay. Captain flying on autopilot using LNAV [lateral navigation] on company stored route everything routine. Captain is a check airman on this aircraft. I am 5-year airline pilot, but only 3 months experience on [glass] equipment. Captain pulled up "hold" page on FMC [flight-management computer] and began to enter data. We were currently descending on rate of descent command to 15,000 ft altitude as assigned. I'm reading holding data to captain as he's entering data via keypad. Both eyes off of primary flight instruments but in altitude capture mode (VNAV) [vertical navigation] and LNAV so both pilots are anticipating automatic level off at 15,000 ft altitude. Captain apparently entered hold at PPOS [present position] instead of at BUCKS in error, and aircraft begins left turn unexpectedly. I, in confusion, not knowing he entered PPOS by mistake, stated, "This isn't right," and saw captain disengage autopilot to stop turn while he goes back to FMC page to find out why the aircraft was turning. I immediately went to the LEGS page to find out where we should be, got to VOR/ILS [very high-frequency omnidirectional radio range/instrument landing system] mode and picked up en route chart to tune in present position. Aircraft descended below 15000'. Horn went off at approximately 14,600 ft. Captain called out "Out of 16 for 15." I said no (looking up) and said we were at 14,500 ft, only cleared to 15,000 ft, that was a low altitude alert, not "1000-ft to go" horn. Immediately pulled back up. No ATC [air-traffic control] communication took place by either ATC or us about event. Within 30 seconds, ATC gave us new frequency and cleared us to 13,000 ft. No conflict or discussion or awareness of event was stated by ATC. Below 15,000 ft. for about 20 seconds.

Cause is obvious—both pilots' attention diverted from aircraft flight path. I suggest hold-page (not used often) be made more user friendly. Have captain tell first officer to watch flight progress while he is correcting other problems. I personally feel that although the FMC is a great tool, it shouldn't be used all the time, especially below FL180. It takes too much pilot attention, especially when newly assigned to the aircraft. Just because the technology exists doesn't mean it should be used. Better pilot training on FMC/pilot problem areas like this should be provided. Also, loss of a flight engineer's eyes is not in the interest of aviation safety. (ASRS#144196)

This incident report illustrates several problems that can occur in the altitude-change task. The crew apparently correctly received the clearance, set the altitude alert, and set up the autopilot to descend and capture the cleared altitude of 15,000 ft. If the crew had done nothing more, the aircraft was set up to level off automatically at the new altitude. Both pilots were distracted from

⁸ The narratives have been lightly edited to improve readability. The numbers that appear at the end of the narratives are ASRS accession numbers.

their routine task of monitoring the altitude change by the nonroutine task of programming the flight management computer (FMC) for the upcoming holding pattern. Normally, no problem would have resulted from not doing the usual checking and monitoring tasks. The pilots might have been out-of-the-loop with respect to altitude, but the autopilot would have performed properly and would have successfully captured the new altitude. The captain's FMS (flight-management system) programming error caused an unexpected left turn. When the captain disconnected the autopilot, the undesired turn was stopped, but disconnecting the autopilot had the unintended and apparently unexpected side effect of cancelling the altitude capture. Both crew members must have been further distracted from monitoring altitude by attempting to diagnose the unexpected autopilot behavior because no mention is made in the report of either pilot making the "thousand-to-go" call at 16,000 ft or seeing or hearing the approaching altitude alert. Diagnostic tasks are notorious for capturing attention from more routine monitoring tasks.⁹ Both pilots in this incident seem to have noticed the low-altitude alert, but the captain apparently interpreted it as the approaching-altitude alert and announced "out of sixteen for fifteen." This erroneous call is a good example of a human behavior with automation phenomena that Wiener has referred to as "primary/secondary task inversion" (Wiener 1985b; 1987). Primary/secondary task inversion describes situations in which a display, in this case the aural approaching-altitude alert, is designed to backup the primary task of monitoring altitude with the primary altitude display the altimeter, but over time the backup display becomes a primary source of information about the aircraft's situation with respect to capturing the new altitude. In this incident, the human monitoring tasks failed and the altitude alerter must be credited with alerting the crew to the altitude deviation.

One way that people detect errors in operating a machine is by noticing that the machine does not respond in the way that they expect it to respond. If the machine does not respond to an input as expected, either the operator has made an input error, the operator's expectation of the machine's behavior to the input is not correct, or the machine has failed. This narrative illustrates the relative effectiveness of two different kinds of unexpected events in alerting the crew to a potential error.

The pilots immediately detected the unexpected left turn. Their expectations were, most likely, that the aircraft would continue on a straight course until reaching BUCKS and only then begin a left turn to enter the holding pattern. Since they had set up the autopilot to capture and hold 15,000 ft they also had an expectation that the plane would continue descending and then level at 15,000 ft. Unfortunately, disconnecting the autopilot to stop the unexpected left turn also had the side effect of disabling the automatic altitude change,¹⁰ and the airplane descended below 15,000 ft. The deviation was not detected until the first officer heard the low-altitude alert.

In the first case, an unexpected event occurred, and the crew immediately detected the problem. In the second, an expected event did not occur, and the crew did not immediately detect the problem. This is probably a result, in part at least, of having an event that is expected to occur, but not immediately. The altitude level-off was expected to happen sometime in the future. Some errors cause unexpected events to occur, and they are the most likely to be rapidly detected. Other errors cause an expected event not to occur, and they are much less likely to be detected when the expected event that is canceled by the error is not expected to occur immediately. The altitude alerter provides a positive event if a level-off at the selected altitude does not occur but in altitude alerters that provide an aural signal only when leaving the target altitude, the visual/aural signal is triggered too late to prevent an altitude deviation.

The reporter of this incident assigned its cause to distraction caused by programming the FMC at low altitude. He suggests less programming at low altitude and more explicit transfer of responsibility to the PNF for monitoring flight progress if the PF is distracted by unexpected problems. Other possibilities for avoiding this particular error are to: (1) delay programming the holding pattern until the aircraft is at the hold altitude, (2) have the PNF do the programming, and (3) have the other pilot check the programming before "executing" the new programmed path. Of course, a simpler hold page or more training on problems and malfunctions for glass-cockpit autopilots might also have prevented this problem.

⁹ The most infamous case is the L-1011 crash in the Everglades (NTSB 1973)

¹⁰ The phrase, "killed the capture," has been used to describe this kind of altitude deviations.

2. Error in Programming Flight-Management System

The following report from the captain of a glass cockpit airliner describes an altitude deviation that occurred during the descent phase of flight when an error was made in programming the FMS.

The first officer was flying the descent with VNAV and RNAV [area navigation] selected. Center cleared us to make a heading change to intercept a radial and cross 30 DME [distance measuring equipment] prior to the VOR at 11,000 ft. I verified that we could make the restriction and observed the first officer program the computer. The new first officer made several errors in the fairly lengthy procedure which required a new waypoint, a radial to the waypoint, crossing restrictions at the waypoint, as well as a turn to intercept the new course. Time became short, and as I attempted to assist him in correcting his errors, I inadvertently inserted the wrong waypoint. (I inserted the radial/altitude rather than the radial/DME as the fix.) I then assisted the first officer in accomplishing the procedure. I did not realize the error until we were leveling at the waypoint, and approach control asked us if Center had given us the crossing restriction.

This error emphasizes that the complex cockpit can, at times, be a hindrance rather than an advantage. A constant verification of data and inputs is required. It would have been much easier to have forgotten the magic of the computer and reverted to VOR navigation while discussing the proper procedures for computer programming after we were on the ground. I feel that below 10,000 ft during the approach phase of the flight, the computer can be detrimental to safety. During this phase, it is essential for one pilot to be looking outside and it is very easy for both heads to be looking at the computer. The navigation display can still be used and is extremely useful for map reference, but probably should be reprogrammed as little as necessary and only when the other pilot is maintaining a vigil for traffic. (ASRS #141870)

In this incident, the autopilot accurately guided the aircraft over the path that the crew had programmed into the FMS but the flight-path description was in error. They created a fix 11 miles prior to the VOR rather than 30 miles prior to the VOR. The reporter is probably correct in stating that it would have been much simpler to abandon the FMS and revert to VOR navigation. Solving their problem with conventional VOR navigation would have involved simply selecting a new VOR radial and doing a little mental arithmetic to figure the appropriate descent rate to meet the crossing restriction. Conventional VOR/DME navigation is characterized by a series of short mental and physical tasks distributed over time. Navigation with an FMS is characterized by a more lengthy initial task of entering the planned route, and then less involvement while the route is being flown. With an FMS, the route has to be completely specified before it can be initiated.

Field studies of cockpit automation (Wiener 1985a; 1989; Curry 1985) have shown that pilots are often reluctant to turn off the automatics and revert to less automatic flight modes. Curry (1985) has recommended that pilots be given “turn-it-off” training. However, the proper solution must await the design of more powerful interfaces to the navigation system or data-link facilities that send the descriptions of flight path in a form that is readable by the FMS, so that pilots do not have to perform the cognitively demanding and error-prone task of translating a clearance from the language of ATC to the language of the FMS.

We will see in other incidents that there is often a mismatch between the pilots’ intentions and the interface that is provided for communicating those intentions to the autopilot and flight management system. In discussing nonaviation computer interface problems, researchers studying the problem of human-computer interaction (Norman 1986; Hutchins et al. 1986) have referred to the problem of translating from a users’ mental intentions to physical actions with the computer as the *gulf of execution*. This incident also provides evidence for a parallel gulf, the *gulf of evaluation*, because the programming error was not evident to the flight crew.

3. A Missing Minus Sign

The next incident report by the first officer of a glass-cockpit aircraft provides a further example of the difficulty pilots can have in translating their intentions to the language of the FMS, and in detecting mismatches between what they intended and what they actually programmed.

On an en-route descent into Dayton our clearance was direct RID VOR, direct DAYTON with a descent to 11,000 ft. The controller gave us a new clearance to cross 10 miles west of RID at 10,000 ft. The captain, being less experienced in using the flight management computer than I, wanted me to show him how to program the descent for the new restrictions. We put the restrictions in the magic box, and for some reason, almost certainly something we did improperly, the machine wanted to make the restriction 10 miles east of RID. By the time we caught the error in the midst of doing checklists and the usual cockpit duties we were too late to make the restriction. Nothing was said and there was no conflict.

There was no question that this mistake was our fault, but it brings up a point. These whiz bang computers and flight management systems are great, but you not only have to watch them like a hawk, they are an error waiting to spring. Flight management systems as they are currently designed do the most and are easiest to use when you need them least—at cruise. The closer you are to the ground, that is on takeoff or approach, the more demanding of attention they are.

Reprogramming for constantly changing clearances, which happens most in the takeoff or approach phases, distracts attention from outside and inside vigilance, and sets up situations where mistakes are likely to occur as the workload increases. To be sure, pilots are responsible for making whatever restrictions there are, not the computers. If we're going to have these sophisticated and really wonderful machines in the cockpit, design effort should be put into making them more useful and less prone to inducing errors in the takeoff and approach phases. (ASRS # 148853)

This and incident No. 2 illustrate an inherent danger in interfaces like the CDU [control display unit] that use a symbolic language to describe the flight path. In incident No. 2, the inadvertent substitution of a symbolic description of the altitude (“11”) for a symbolic description of the distance (“030”) of the fix from the VOR resulted in a 19-mile error in the fix location. In this incident, an error in inputting a single character (“RID/10” versus “RID/-10”) resulted in a 20-mile error in the location of an altitude-crossing restriction. The power of symbolic systems is that small descriptions can specify large processes. The danger is that small mistakes in specification can lead to large errors. The incommensurability of the magnitude of causes and effects in symbolic systems can create situations that take operators by surprise.

This narrative also illustrates another factor that was observed in several reports: in-flight training, in which one pilot attempts to instruct the other pilot on the use of automatic equipment, can distract both pilots from the primary task of maintaining altitude awareness. In-flight training on the operation of complex systems such as the autoflight and FMSs is going to happen, but it should be done during the low-workload cruise phase of flight. Unfortunately, as the report of this incident points out, the clearances that require complex reprogramming usually occur during the already busy climb and descent phases of flight.

4. The Green Arc

The following incident report from the captain of a glass-cockpit airliner describes an altitude deviation in which unexpected behavior of the green altitude arc on the map display helped the captain detect the problem. The green arc is a symbol on electronic map displays that indicates where the aircraft will be on the map when it reaches the altitude set in the altitude alerter window if the aircraft maintains its current rate of descent. It provides the pilots with a tool for monitoring progress toward the achievement of a crossing restriction without requiring FMS programming.

Approaching Denver over NW arrival gate, I programmed the FMC for a VNAV destination to cross DRAKO at FL200, 250 KIAS and Denver at 11,000 ft. Approaching DRAKO, Denver approach controller cleared us to cross the Denver

313° radial at 10 DME fix at 13,000 ft due to turbulence. I disengaged the autopilot and hand flew the descent using the green altitude arc over the 10 DME fix waypoint inserted by the first officer. This arc is predicated on the altitude selected in the altitude window of the MCP [mode control panel]. As we progressed, I noticed I was falling behind the required descent profile. Checking the altitude window, I saw 11,000 ft. I asked first officer if we had been cleared to 11,000 ft, thus cancelling the 13,000 ft restriction. He replied that the clearance was to cross the fix at 11,000 ft. Apparently when he had inserted the waypoint, he had put in 11,000 ft instead of 13,000' and later set the MCP altitude window to agree with it. Anyway, I increased our descent rate only to have approach control point out our error as we neared 11,000 ft. Recleared to 11,000 ft, we completed the flight.

Due to the turbulence and the fact that we normally are at 11,000 ft at that point in the profile, I failed to have the first officer verify our clearance before descending below 13000'. We were in the tenth hour of a duty day that began early a.m. and fatigue may well have been a factor. Further, I have resolved to do less VNAV techniques in the approach phase. EFIS (electronic flight instrument system) is great stuff, but there is a tendency to try to do too much with it too long. I would be interested to know if you can detect a trend of increasing incidents concerned with glass cockpits. (ASRS #144996)

This incident further illustrates the role that expectations can play in detecting errors or, in this case, almost detecting errors. The captain noticed that the aircraft was “falling behind the required descent profile.” This statement implies an expectation by the captain of how the aircraft should have been proceeding. The unexpected behavior triggered a check of the altitude set in the altitude-alterer window. This value, 11,000 ft, was different than the captain expected and he queried the FO (first officer). Here he almost caught the error that caused this altitude deviation. As he mentioned later in the report, “I failed to have the first officer verify our clearance before descending below 13,000 ft.” Another expectation that the captain had that influenced his decision not to check the clearance with ATC was that “we normally are at 11,000 ft at that point in the profile.”

Similar to comments made in several other reports in our sample, the reporter mentioned that fatigue may have been a contributing factor. A number of the reports concluded with a phrase similar to “and we were tired.” This is important because the redundant checking that enables the system to be error tolerant are a sort of “luxury” task. If one has the time to do them, they are done, but when work load is high and the pilots are fatigued, redundant checking tasks are often curtailed first, when they are needed most.

5. An Unexpected Clearance and a Tired Crew

The following incident reported by the first officer of a glass-cockpit aircraft further illustrates the way expectations and fatigue can affect task performance.

Atlanta tower will clear us to 4,000 ft on departure and give us our first turn and always immediately hand us off to departure control, who will always clear us to 10,000 ft (usually long before we get to 4,000 ft). What happened: On departure, Atlanta tower gave us our first turn and that was all. The captain and I were so conditioned to the continual climb out from Atlanta that we continued through 4,000 to 4,800 ft before I discovered it and called tower for a hand-off while the captain stopped the climb and began a descent back to 4,000 ft.

This was the second day of a three-day trip and the captain and I had had a short night before (minimum legal rest). I freely admit that often I feel exhausted after and during my trips. This condition often accounts for mistakes and oversights not related to the safety of the flight. However, I've got to consider the possibility that the chances go up for making such a mistake that does affect the safety of the flight in that condition. Consider the minimum 8-hour rest. Off duty at 10:00 P.M. Hotel van shows up at 10:15 P.M., arrive at the hotel at 10:30

P.M., in bed by 11:00 P.M., wake up call at 04:45 A.M., leave hotel 05:45 A.M., airport at 6:00 A.M. Net 5 3/4 hours sleep. Get real guys! The FARs allow this but nobody seems to acknowledge that it's unreasonable. Not the Feds, not the company, and even many pilots have just given up complaining. Why? (ASRS #154424)

Here the crew flew the departure that they expected to fly, not the actual clearance. In the human error literature, this is called a capture error (Norman 1981; Reason and Mycielska 1982; Reason 1990). A common example is intending to run an errand on the way home from work and finding yourself at home with the errand not done. This phenomenon is more prevalent when complicated by crew fatigue.

This incident is another example of the difficulty in detecting the absence of an expected event—in this case the absent hand-off and clearance to 10,000 ft.

6. One-Pilot-off-Radio Incident

This incident was reported by the first officer of a conventional-technology aircraft. It illustrates a common breakdown in the error-tolerant properties of the cockpit system in both the conventional and glass cockpits. One of the pilots was off the ATC radio frequency, talking to his airline operations or getting the ATIS (automated terminal information system) airport information, and did not detect a clearance communication error by the other pilot. The reporter suggests that this class of incident might be reduced when data-link is utilized and pilots can receive ATIS information and company communications without leaving the traffic-control frequency.

While on a flight from EWR to SYR, we began our approach into the SYR area. I was first officer and the captain was flying the aircraft. I left Center frequency in order to listen to the ATIS and talk to company operations. When I returned to Center frequency, the captain informed me that we had been cleared to cross 35 n. mi. east of SYR at 10,000 ft. After descending to 10,000, we were instructed to contact SYR approach. The approach controller was very surprised that we were at 10,000 ft. He told us that the normal procedure was . . . to clear inbound aircraft down to 11,000 ft and then instruct them to contact SYR. This information of course caused us to question whether or not the captain had correctly understood the clearance. The SYR controller advised us that no conflict had resulted.

This is yet another case of a possible clearance deviation while one flight crew member is not monitoring the ATC frequency. All ATC communications should be monitored and verified by each pilot whenever possible. This is especially important in the case of altitude assignments and crossing restrictions. Situations such as this could easily be avoided in ACARS-equipped aircraft by using ACARS [ARINC Communications and Reporting System] to obtain ATIS, gate, and other information. The technology required to do this has existed for years. Unfortunately, it will probably take a major accident before the airlines begin to fully utilize this equipment. (ASRS #147683)

As we will discuss in the next section on the properties of error-tolerant systems, when neither pilot is monitoring the ATC frequency many potential opportunities for detecting and correcting communication errors are lost.

7. Lack of an Aural Approaching-Altitude Alert Blamed

In earlier reports on altitude deviations (Factors . . . 1977; Human . . . 1978), several pilots complained about the aural warn that sounded to alert them that the aircraft was approaching the altitude set in the altitude-alerter window. Pilots complained that this chime went off so often that they became habituated to it and often did not hear it. It was also argued that because of the presence of this aural reminder, pilots often did not monitor altitude and relied solely on the aural alert for awareness of the approaching level-off. Wiener (1985b; 1987) described this change as an example of primary/secondary task inversion. Because of these arguments, the aural part of the

altitude alert was made optional and is only a feature in some glass-cockpit airliners.¹¹ Our sample of reports contained several cases like that noted below in which the reporter claimed that if there had been an aural alert, the reported altitude deviation would not have occurred.

While on vectors for an ILS runway 13L at JFK, we were instructed to descend and maintain 3,000 ft. We vacated 7,000 ft and began in-range checklist. Aircraft was being flown on autopilot (autopilot is equipped with altitude capture). As I was completing the approach briefing, approach calmly advised us to climb and maintain 3,000 ft. (No sense of urgency in his voice). We had allowed the aircraft to descend to 2,200 ft. I immediately disconnected autopilot and climbed back to 3,000 ft. (The autopilot indicated an altitude capture, but it continued to descend). The flight continued without further occurrence.

The controller never mentioned anything more about the altitude excursion. I don't think that our altitude excursion resulted in any traffic conflicts. This occurrence could have easily been avoided by monitoring the aircraft's progress throughout the descent. In my experience in this aircraft type, I had never witnessed it not capture an altitude. This good reliability of the autopilot may have left some complacency within me. The aircraft isn't equipped with an aural altitude alert, but rather with a dimly lit light. This occurrence wouldn't be a problem in the future with proper altitude awareness and constant monitoring of the aircraft. I believe that we wouldn't have busted our altitude by 800 ft if our aircraft was equipped with an aural warning. (ASRS #147164)

Always having or never having an aural alert are not the only options available for the altitude-alert system. Lyddane (1991) has argued that the aural altitude alert should only sound if the altitude-alerter system detects that an immediate level-off is required to capture the altitude set in the altitude window. Pilots should not habituate to this alert because it should only sound in the rare situations when an altitude deviation is about to occur. It also should alert the pilot so that action can be taken in time to prevent the altitude deviation.

This report is an example of crew complacency with the automatic system. As the reporter mentions, both pilots were running a checklist and conducting the approach briefing and not performing the normal checking of aircraft altitude. The reporter states that "This occurrence could have easily been avoided by monitoring the aircraft's progress throughout the descent." Complacency and distraction by apparently more important primary tasks appear to be the main reasons why redundant monitoring and checking tasks are not accomplished, yet it is the monitoring and checking tasks that give the system its error-tolerant properties.

8. High Performance Low Altitude Level Off

The new-technology aircraft differ from older aircraft in several ways other than the level of automation in the cockpit. The goal of having good single-engine performance when one engine is shut down and the aircraft is heavily loaded has resulted in remarkable climb performance in a lightly loaded aircraft with both engines operating normally. The following incident illustrates how lack of pilot familiarity with this high-performance aircraft and a low-altitude level-off can result in a speed and altitude deviation.

Taking off from ORD in a [large transport] with a light load and maximum take-off power (engine anti-ice on). The first officer, just out of training, was flying the leg while I handled communications. Each of us had only done one previous leg in a [large transport]. (I have several hundred hours in [wide body transport].) The combination of cold weather, maximum power, and a nearly empty aircraft caused the airspeed to increase extremely rapidly after liftoff. The first officer was reluctant to raise the nose to the extreme angle required to maintain 250 (in this case, probably better than 25°). When I saw the airspeed zipping though 270, I warned him to slow down and he disconnected the

¹¹ See FAR Paragraph 91.219.

autothrottles, manually retarding power and raising the nose just as the flight director went to altitude capture (between 3,500 and 4,000 ft). Attempting to level at 5,000 ft, we overshot by 200-300 ft (still fast), when we were cleared to 14,000 ft. (I don't really know whether we actually broke 5,300 ft before being cleared up.) I punched flight-level change, but the autothrottles refused to engage initially. In the confusion over exactly what was wrong, we both were slow to respond to several heading changes, which understandably annoyed the controller.

Nothing really serious here, except the same old story. Both of us were engrossed in trying to figure out why this computerized marvel was doing what it was, rather than turning everything off and manually flying (which we finally did) until we could sort things out. This is a common tendency in this type of cockpit, but our unfamiliarity with the super high performance of the LGT [large transport] was a contributing factor. It really is a handful to takeoff and level at a low altitude and seems to require an almost immediate power reduction to maintain a reasonable nose attitude at low weights. (ASRS #134179)

In addition to the problems caused by the crew's lack of familiarity with the high performance of their twin-engine aircraft, this incident illustrated two other problems. First, like the crew attempting to set up the hold at waypoint BUCKS 9 (incident No. 1), this crew mentions being distracted by attempting to determine why the autoflight system was not performing as they expected. Second, as reported in prior examples, this crew appears to have difficulty translating the departure clearance from ATC language to the language of the autoflight and FMS. In this incident, the pilots would probably have liked to have been able to instruct the autoflight system to adjust the power after takeoff so that the air speed would not exceed 250 knots and the pitch angle would not exceed some upper limit. It is not possible to enter a combination of an altitude target and pitch-attitude constraints with the autopilot MCP or FMS in today's cockpits. In these situations, pilots could probably more easily achieve their objectives by performing a reduced-thrust takeoff or by not using the autothrottles at all and controlling power manually instead.

9. Ambiguous Shorthand and a Misset Altimeter

The following incident illustrates a classic cause of an altitude deviation that is possible on both conventional- and glass-cockpit aircraft. Aircraft altimeters depend on barometric pressure to measure altitude. If the reference barometric reading is not entered correctly, the altimeters will display the wrong altitude and, unless detected by the pilots, the aircraft will level at what appears in the cockpit to be the correct altitude but is not, because the altimeter's barometric reference has been set incorrectly.

On descent, captain misread ATIS altimeter setting (placed wind in barometer window as opposed to altimeter setting); i.e., 3009 instead of 29.57. Autopilot leveled aircraft at 10,000 ft, but aircraft was actually at approximately 9450'. Controller asked our altitude and gave the current altimeter setting, at which time the captain reset his altimeter and the aircraft returned to proper altitude.

I have noticed on several occasions (in a two-man crew) that in busy environments such as ORD, cross checks on altimeters and other items are not performed as quickly as they should be. There were two approaches in use, but none had been specified for use for us, so both needed to be studied, along with steering for vectors, and complying with ATC instructions and performing other cockpit duties. I had written the ATIS information on our company changeover information paper and presented it in the normal fashion: 80 3H 3009 29.57 ILS 27R ILS 27L.¹² We were both busy, as is normal at this phase of flight--especially

¹² This 32-character string—"80 3H 3009 29.57 ILS 27R ILS 27L"—decodes to the following: The ceiling is 8,000 ft, visibility is 3 miles in haze, wind direction is 300°, wind speed is 9 knots, barometric pressure is 29.57, expect an instrument landing system (ILS) approach to either runway 27L or 27R.

on a 2-man crew, and captain misread and misset the altimeter and I failed to cross check it in a timely fashion to prevent an altitude deviation. (ASRS #133242)

Cryptic abbreviations of aviation information combined with wind descriptions that make plausible barometric settings encourage this kind of error. This is another example in which a small error in processing symbolic information can lead to a large error in the altitude of the aircraft—in this case an error of 550 ft.

Again, the introduction of data-link facilities that allow pilots to receive ATIS information when time permits should reduce errors of this kind. Of course, data-link technology will do little for this problem as long as the format of the presentation of weather data continues to incorporate these abbreviations. For example, many glass-cockpit airplanes currently receive destination weather reports via data-link. Here is an actual example of a destination weather report received via ACARS on a recent flight to San Francisco:

```
SFO 1450 CLR 8 193/  
46/41/2303/010/FEW  
SC CI E / FH ALQDS/  
103 1501 / SOSF 50/  
1206
```

Looks like a nice day, doesn't it? The technology exists to translate this into a format that would be much easier to read. This also contains information that is of interest to meteorologists, but not to pilots. For example, the "193" in the first line means that the sea level barometric pressure is 1019.3 millibars, information that is of no use to pilots. The barometer setting they need for their altimeters is in the second line, "010", meaning the local barometer setting is 30.10 in. Hg. Many pilots have no idea what those first digits encode, and there is no reason they need to know.

10. Misplaced Trust in the Autoflight System

In the following narrative, the pilot provides a good summary of how the lack of checking and cross-checking tasks during flight with the autoflight system engaged can lead to an altitude deviation.

Another altitude bust in the "Brand X" aircraft. I've never in my career had so many altitude busts until I got in this airplane. The scenario is all too familiar. One pilot involved in approach miscellany (ATIS, checklists, etc.) the other pilot trusting the autopilot to make the level off and getting involved in something else at the last minute. Next is the sound of the altitude warning and a mad scramble to level the aircraft off before ATC notices. This time we made it down to 500 ft too low due to a high rate of descent engendered by a crossing restriction given too close in by ATC. My fault for assuming the first officer had it under control. His fault for assuming the autopilot had it. The designers/flight managers/flight instructors fault for encouraging maximum utilization for the autopilot which has the net effect of taking the pilot out of the loop in spite of your best intentions. Human beings are not good "monitors" of machinery; the mind tends to wander if there is no need to actively control the process. (ASRS #144385).

This report highlights the fact that once the automatic systems are programmed, supervisory control tasks remain for the pilots.

FIELD STUDY FINDINGS

The altitude-deviation incidents described in our sample and the pilot/reporters statements concerning the probable causes of those incidents are in general agreement with the findings of an extensive field study of pilots of a glass-cockpit airliner performed by Wiener (1989). He reported that “Many pilots, while reporting that they enjoyed flying a modern plane, also indicated strong reservations in two critical areas: (1) safety, and (2) workload reduction. As for safety, many of the crews expressed the view that automation¹³ may have gone too far, that they felt they were often ‘out of the loop,’ probably meaning that they tended to lose situation awareness, and that they feared that automation led to complacency.” With respect to workload, there was strong disagreement, but at least half the respondents reported concern that automation actually increased workload, that workload was increased during phases of flight already characterized by high workload, and decreased during periods of low workload. Many crews reported that in times of heavy workload, they tended to “click it off,” that is to revert to manual modes of flight guidance, because they did not have time to do the programming necessary to exploit the automation. Wiener refers to this last finding as the “paradox-of-automation”; i.e., when the workload is highest, it is often not used. Wiener attributes this problem both to the difficult-to-use interface with the FMS and to the many ATC clearance changes that are common in busy terminal areas.

BUILDING A SYSTEM DESCRIPTION

The altitude-change task consists of several components. Some of these can be specified a priori from knowledge of the task. Figures 2 and 3 were constructed in this way. Other elements may be specified in the various documents that are involved in training pilots, regulating their actions, or promoting a “pilot culture.” For example, the Airman’s Information Manual (AIM) specifies pilot responsibilities upon issuance of a clearance as follows:

- a. RECORD ATC CLEARANCE – When conducting an IFR operation, make a written record of your clearance... .
- b. ATC CLEARANCE/INSTRUCTION READBACK – Pilots of airborne aircraft should read back those parts of ATC clearances and instructions containing altitude assignments or vectors as a means of mutual verification. The read-back of the “numbers” serves as a double check between pilots and controllers and reduces the kinds of communications errors that occur when a number is either “misheard” or is incorrect... .
- c. It is the responsibility of the pilot to accept or refuse the clearance issued. (AIM ¶ 265).

According to the AIM, upon receiving a clearance, the crew should record the clearance, decide whether the airplane can comply with the clearance, and then either read it back if it is accepted or notify ATC that the clearance cannot be accepted. That is a normative view of clearance handling. Observations of actual line operations and full-mission simulation studies show that some elements are sometimes omitted. Clearances are usually written down when they are complex, for example when they involve a hold, or a revised departure procedure. Clearances that consist of simple altitude or heading change instructions are usually not written down. Thus, the actions that are actually performed by any particular crew in any particular circumstance will usually be a subset of a larger set of possible actions. In order to understand the properties of the information processing system that responds to the ATC clearance, we would like to know the full range of possible crew actions.

Each of the 10 ASRS narratives presented above refers to some of the component of the altitude-change task and leaves others unspecified. At least one of the narratives refers to each of the different elements given by the AIM, although no single narrative refers to all three. Those elements that are unspecified may have been omitted because they did not occur or because the reporter simply did not find them essential to the point being made. By aggregating across the reports or superimposing the elements of each on a common framework, we can create a more

¹³ The term Automation in Wiener’s report (as in ours) refers to the flight guidance automation, the autopilot, autothrottle, and the FMS, not the aircraft systems or the warning system automation.

complete description of the task. This will not be a description of any particular performance of the task, but will be a superset of the actions that actually occur during any particular altitude change. It will be a general framework that shows the possible components of a task.

Terms of Task Analysis

Information relevant to the altitude-change task is represented in several different ways, and we can track its movement through the cockpit system. A clearance to a new altitude enters the cockpit as a sequence of spoken words that are heard by the PNF and read back, often slightly transformed, to the controller. The specification of the altitude is then transformed by the PNF into a setting on the altitude alerter. Notice that the altitude specification is in some sense the “same” information whether it is represented in spoken words, as a string of written characters, or in the digits visible in the display window of the altitude alerter. Speech and writing and the display window are three different *media* in which task-relevant information may be represented. A medium is said to represent some particular piece of information by virtue of having a particular physical state. For example, the altitude-alerter display window represents the altitude thirty three thousand feet when the digits “33000” appear in order in the window. The crew moves information through the cockpit system by translating the representation of information in one medium to a representation in another. We will characterize the information processing of the cockpit system as a whole in terms of the movement of task-relevant information by agents from one medium to another. Each medium represents information in a unique way.

Each medium in which information is represented has properties that determine how long the information expressed in that medium lasts, and how accessible it may be to the members of the crew. Speech, for example, is usually highly accessible, but it does not endure in time, so information that is expressed in that medium must be attended to at the time it appears. Notes written by a pilot on a side window clipboard are less accessible to that pilot than speech is, and may be completely inaccessible to the other pilot. However, written records endure in time, so that information that is expressed in that medium lasts and may be referred to many times. These may seem to be trivial observations, but the cockpit information processing system is made up of these media, and little facts like these determine the information processing properties and error-tolerant properties of the system as a whole.

System of Representations

The set of representations involved in flight path management is quite complex. Much of it, however, is directly available to an observer in the cockpit, and each of the ASRS narratives quoted earlier refers to aspects of it. In the following subsection, we develop a framework for the altitude-change task that includes the elements of the task mentioned in the narratives. This composite picture will not be a description of any particular altitude change event. It will be instead a description of the possible paths that task-relevant information can take in the cockpit system. Each path of information through the cockpit is made up of a number of steps in which the information is moved from one medium to another. Most of the transformations of information are accomplished by members of the crew. We will try to show that the error-tolerant properties of the system as a whole depend on the organization of these possible trajectories, as well as on the cognitive properties of the individual crew members. In particular, we will show that errors are often detected as mismatches between representations that are expected to be congruent.¹⁴

Theory of Action

The widespread use of personal computers for many tasks has resulted in renewed interest in how people interact with computers and how they should interact with them. In an attempt to develop a better understanding of how people use computers to accomplish tasks, Norman and his

¹⁴ Representations may be congruent without being identical as long as they have the same interpretation. That is, at the level of meaning, the clearance “slow to one nine zero knots” and the read back, “pulling it back to one ninety” are congruent, although at the level of expectations in auditory memory they are discrepant.

colleagues, (Norman 1986; Hutchins et al. 1986) have developed a *theory of action* that attempts to specify the mental activities a user must execute to perform a task with a computer. It is one of the few descriptive models of human behavior that recognizes the importance of evaluative feedback tasks.¹⁵ Their model is concerned with how a person bridges the gap between task goals expressed in psychological terms and the machine's mechanisms and states that are expressed in physical terms. They refer to the discrepancies between the psychological goals and the physical variables as two gulfs that must be bridged by the system: the gulf of execution and the gulf of evaluation.

Figure 4 shows the seven stages of user activity that are involved in performing and evaluating a simple task on a personal computer. To see what this means in the context of the altitude-change task, consider setting the altitude alerter. The first stage is to establish the goal. For this task, that would be something like deciding to set the alerter display to the cleared altitude. The next three stages, forming the intention, specifying the action sequence, and executing the action, are required to translate from the person's goal expressed in psychological terms to a sequence of inputs that the machine understands—the machine's input language. The pilot intends to produce a particular reading in the display. The action specification in this case would involve grasping and rotating the setting knob of the altitude alerter. Prior knowledge that clockwise rotation produces larger values might be involved in the formation of the action specification, or it might be rediscovered in action by turning the knob and monitoring the result. The action specification is then executed as action in the world. This produces a new physical state in the device.

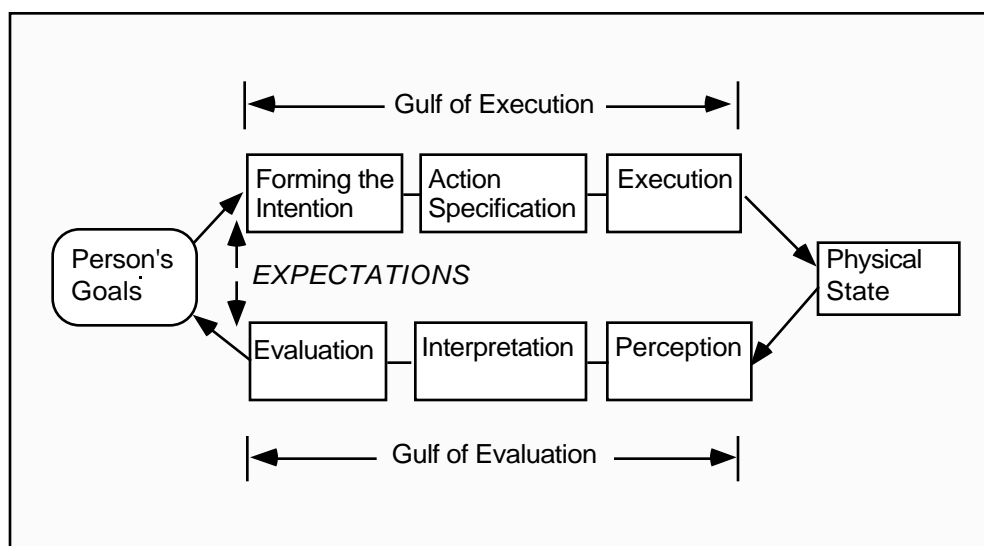


Figure 4. Theory of action.

These three cognitive tasks—intention, specification, and execution—constitute the gulf of execution. The cognitive processes involved in performing these elements bridge the gulf of execution. The kinds of processes and the amount of cognitive resources required to bridge the gulf of execution may depend both on the task demands of the environment and the knowledge of the operator. The gulf could be narrowed either by changing what the operator knows or by changing the design of the devices with which the task is performed.

Once a new state has been produced, one must determine whether it is the desired physical state. Is the number in the altitude display a representation of the cleared altitude? To determine this, the state must be perceived, the perception interpreted as a number, and the resulting representation must be compared with the intended state. These feedback elements constitute the gulf of

¹⁵ An exception to this statement is the many models developed in the field of manual control theory (Baron et al. 1990).

evaluation, and the processes involved in performing these elements bridge the gulf of evaluation. In the case of setting the altitude alert, bridging the gulf of evaluation is relatively simple. If the goal was to produce a particular kind of flight path with crossing restrictions by programming the FMS, the execution and evaluation process would have been much more complex. In that case, the high-level goal of causing the airplane to perform a particular kind of climb would subsume other goals, each of which might spawn a cycle of action and evaluation. Again, the gulf of evaluation can be narrowed either by changing what the task performer knows or by changing the design of the devices with which the task is performed. The design of the human-machine interface plays an important role in determining how much cognitive work the user must do in bridging the gulfs of execution and evaluation.

Acting may involve both the production of action and an evaluation of the extent to which the achieved outcome satisfies the actor's intentions. Mismatches between intentions and outcomes are evidence of error and the feedback inherent in the theory of action provides opportunities for operator detection of errors.

Chronology of Altitude-Change Task Elements

There are at least seven distinguishable elements of the altitude-change task.

1. Hear and record the clearance. Hearing the clearance was not mentioned in the AIM description, but it is an obvious precursor of the other elements. Typically, if a clearance is to be recorded, it will be written down as the clearance is spoken. We have seen that this is optional for simple clearances. When a clearance is written down, it may be recorded before any effort is made to understand what it implies for the airplane.
2. Understand the clearance and decide whether the airplane can comply. Once the clearance is heard and understood, the crew must decide whether they will accept it.
3. Read back the clearance or notify ATC that the clearance cannot be accepted. If the clearance is accepted, then the crew must form a high-level goal to comply with the clearance. This goal is realized through the satisfaction of a number of sub-goals. In glass-cockpit autoflight operations, the following sub-goals must be satisfied.
4. Create a machine-readable description of the flight path implied by the clearance. Entering the altitude in the altitude alerter is usually part of this task but not always all of it.
5. Select a mode of operation for the autoflight computers that uses the appropriate parts of the flight path description.
6. Couple the autoflight computer outputs to the flight controls.
7. Monitor the performance of the autoflight system and intervene if necessary.

The cognitive activities involved in the accomplishment of each of these subgoals represent a loop around the seven stages of action. In addition, the satisfaction of the higher-level goal of complying with the clearance must be evaluated. We may therefore also consider all the actions described above as elements of the performance of the task of causing the airplane to comply with the clearance, and we may consider the monitoring activity that follows these activities (sometimes for many minutes) as the evaluation of the adequacy of the actions performed.

If there were only one way to describe the flight path, the altitude-change task would probably be a relatively simple one. However, the modern flight deck contains several forms of description of desired aircraft behavior. Since the language in which flight path is described to the machinery is quite different from the language in which the clearances are communicated between pilots and controllers, it is sometimes difficult to determine whether the description passed to the autoflight systems will produce the results described by the clearance. The two most apparent classes of machine-readable descriptions of flight path are those entered in the MCP¹⁶ and those entered in the FMS. The MCP actually provides for the creation of partial descriptions of flight path and for the selection of the mode of operation that links parts of the descriptions to the autopilot.

¹⁶ "Mode control panel" is Boeing terminology; Douglas Aircraft Company calls the same device the flight guidance control panel and Airbus calls it the flight control unit..

Let us now consider in some detail the cognitive processes involved in the performance of each element of the altitude-change task.

Hear and Record the Clearance

Each report narrative (except No. 8) begins with a reference to a spoken clearance. The clearance is a specification of some aspect of the flight path of the aircraft. The initial step in the trajectory of the clearance through the cockpit system is from the medium of speech to the medium of the pilots' mental representations of the meaning of the clearance. Pilots form at least two kinds of internal representations of the clearance: auditory and semantic. The auditory representation is simply a representation of the sounds of the spoken words. The semantic representation is a representation of what the clearance means in terms of the flight path of the airplane. Pilots may have expectations with respect to either sort of representation and expectations provide the opportunity for checking the accuracy of the representations (see Fig. 5).

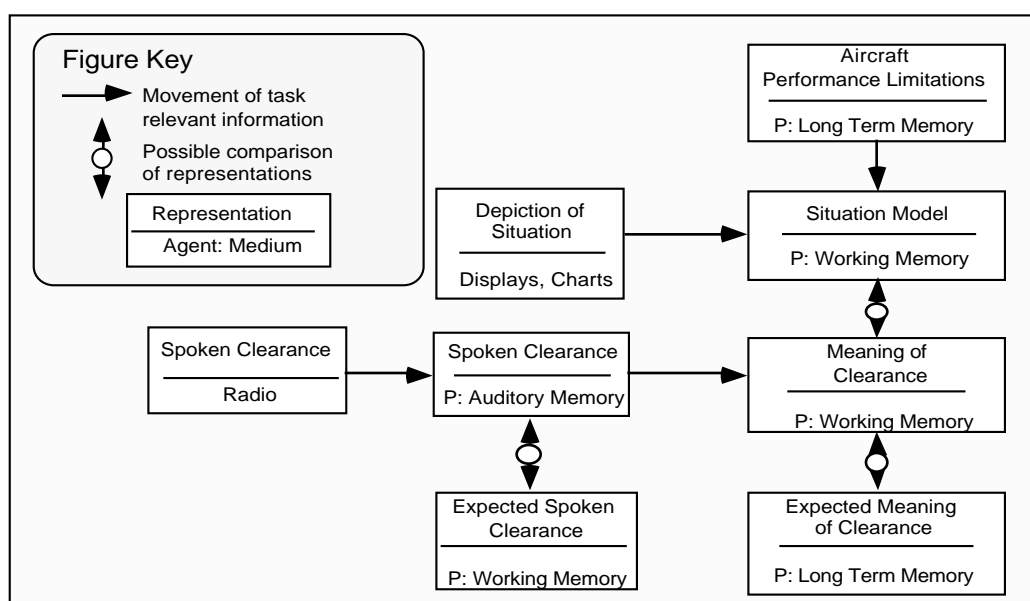


Figure 5. Information flow and checks in checking the clearance.

Perception is driven in part by expectations. Expectations permit crews to “hear” and reconstruct spoken clearance information that is degraded by noise. Unfortunately, often to a large extent we see (or hear) what we expect to see (or hear). Expectations permit pilots to reconstruct the wrong information when expectations do not match the clearance that was sent, or may result in pilots not hearing that which is not expected (Hutchins and Klausen 1992).

Pilots also sometimes write the clearance on a piece of paper. Incident No. 1 seems to involve this representation, since the FO makes reference to reading the clearance information to the captain. Creating the written representation from the spoken representation involves another translation from one medium to another and from one set of representations to another. This translation is relatively automatic for those literate in the language. Because of frequency congestion, clearances are often spoken very rapidly, and a shorthand is required to record them in real time. Incident No. 9 demonstrates some of the problems with abbreviated written representations. The four-digit string “3009” is a reasonable representation of either barometric pressure in inches of mercury (30.09 in. Hg.) or of wind direction and speed (300° and 9 knots). This ambiguity led to a misinterpretation and an altitude deviation.

Understand the Clearance and Decide on Acceptance.

The meaning of the clearance is processed by the crew with respect to an existing model of the aircraft situation and possibly to a memory of clearances previously received in similar circumstances. These situation models may entail complex expectations about the clearance to be received at any point in flight. This is stated most clearly in narrative No. 5.

Atlanta tower will clear us to 4,000 ft on departure and give us our first turn and always immediately hand us off to departure control, who will always clear us to 10,000 ft (usually long before we get to 4,000 ft.

The good news is that expectations of this sort may permit pilots to detect errors made by controllers. A mismatch between the meaning of the clearance received and pilot expectations may direct the pilots' attention to the clearance. The bad news is that action as well as perception may be expectation driven, and if expectations about flight path do not match the cleared flight path, a crew may do the wrong thing. The very strong expectation cited above apparently led the crew to begin to fly their normal departure procedure in the absence of a clearance to climb above 4,000 ft.

The expectations can also be complex and go beyond anticipating particular clearance information at particular times. Narrative No. 1 shows that the pilot had a model of the situation, not only of his own airplane, but of the surrounding airspace, of the problems faced by ATC, and of the sort of solution ATC would seek to a weather-induced problem:

Descending from higher flight levels to 15,000 ft on Center clearance had anticipated and received clearance to hold at BUCKS due to anticipated weather delay.

Upon hearing the clearance, the pilots should determine whether the airplane is able to comply with the clearance. Such determination interacts with the current situation assessment of the crew and with the performance limitations of the airplane, especially with respect to altitude changes. Can the airplane climb or descend quickly enough to make a specified crossing restriction? Will it have an adequate buffet margin at the assigned altitude at current gross weight? Will this altitude and heading provide adequate terrain clearance? Answering these questions may require the pilots to make use of other information in the cockpit. The pilots' understanding of the aircraft situation against which the ability to comply with the clearance is tested may not be entirely internally represented, but may be constructed in interaction with a chart or approach plate, or other depiction of the physical space around the airplane.

The ease with which these questions can be answered will depend, at least in part, on the way the relevant information is represented. If the crew has a moving-map display, the consequences of changes in lateral guidance may be relatively easy to evaluate. In current technology, questions about changes in vertical guidance are more difficult to answer because there is no graphic display of the vertical flight plan. We will return to this issue in the discussion of recommendations. The degree to which crews actually evaluate the consequences of complying with clearances seems to be highly variable. Although this assessment should be part of every clearance handling event, it is explicitly mentioned only in narrative No. 2: "I verified that we could make the restriction...".

Since the clearance may be heard by both pilots, all the processes involved in checking the clearance may take place twice, once in each pilot. This potential for redundant processing by the two pilots is one of the most important properties of the cockpit information processing system. At least three opportunities for the detection of error through the comparison of representations exist within each pilot.

Read Back the Clearance

If the clearance was written down, the read-back usually consists of literally reading the written record. If the clearance was not written down, the read-back must proceed from memory. This can be difficult if the clearance included several elements such as an altitude, a heading, and a frequency to use in contacting the next ATC facility. There are many opportunities for error in this procedure, because altitudes, headings, ATC frequencies, and aircraft call signs all consist of strings of numbers. If it is decided that the airplane can comply with the clearance, the portions of the clearance specifying altitudes or headings should be read back to the controller. The read-back

provides an opportunity for redundant checking on the clearance as heard by the pilots. It gives the controller an opportunity to compare what was told to the pilots with what the pilots read back (see Fig. 6). Two of the reports refer to possible failures of the read-back to detect an error. In narrative No. 4, the captain reports that when the airplane performance indicated a steeper than expected descent:

I asked first officer if we had been cleared to 11,000 ft, thus cancelling the 13,000 ft restriction. He replied that the clearance was to cross the fix at 11,000 ft.

In this case, the captain is using the first officer's memory as a check on his own memory for the clearance received. This crew might also have asked ATC to confirm the clearance, but unfortunately, they did not.

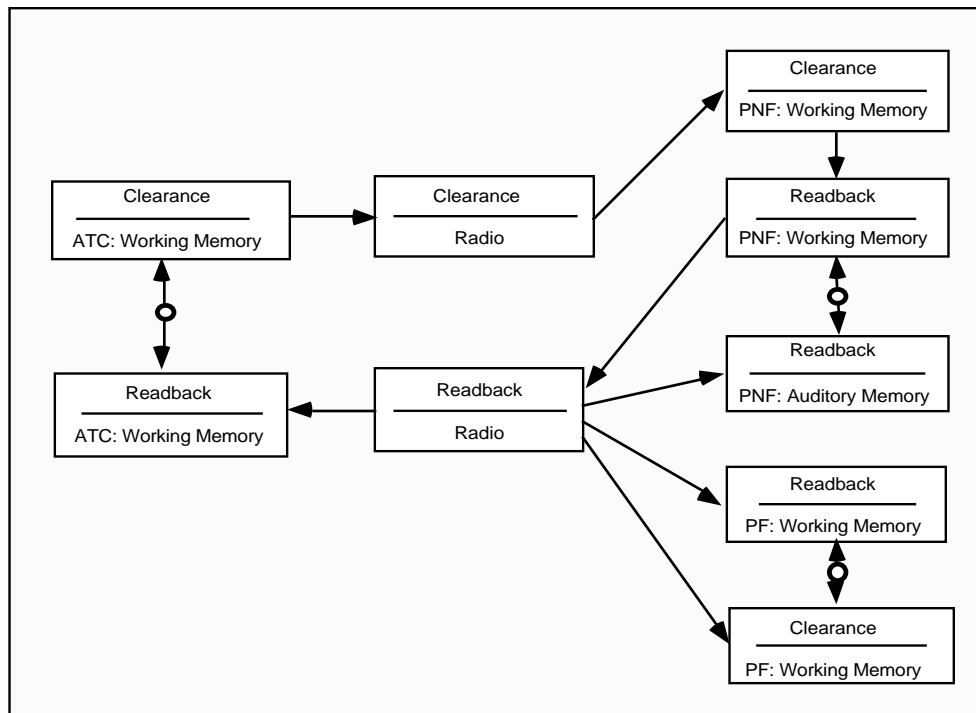


Figure 6. Information flow and checks involved with the read-back check.

The redundant checking that is provided by both crew members usually monitoring ATC communications may break down when one of them leaves the frequency to attend to other work. Three of the incident reports refer to this loss of redundant cross-checking. In narrative No. 6, one of the crew was off-frequency getting ATIS when a clearance was received:

When I returned to center frequency, the captain informed me that we had been cleared to cross 35 n mi east of SYR at 10,000 ft. After descending to 10,000, we were instructed to contact SYR approach. The approach controller was very surprised that we were at 10,000 ft. He told us that the normal procedure was...to clear inbound aircraft down to 11,000 ft and then instruct them to contact SYR. This information, of course, caused us to question whether the captain had correctly understood the clearance... . This is yet another case of a possible clearance deviation while one flight crew member is not monitoring the ATC frequency. All ATC communications should be monitored and verified by each pilot whenever possible.

The read-back also permits checking within the cockpit itself, because as one pilot reads the clearance back, the other pilot, who usually has also heard the clearance, may be able to detect an error in the read-back. If both crew members are on frequency when clearances are received, there is also a possibility of redundant checking of the read-back by the other member of the aircrew (Hutchins and Klausen 1991).

So far, we have seen that the media in which information about the flight path are represented include speech (the clearance, the read-back), the mental representations of aircraft situation in the minds of the two pilots, and the charts and other graphical depiction of the flight environment. Notice that the different media support descriptions of flight path in different languages. That is, the clearance is in spoken aviationese, whereas the pilots' mental representations may be in some form of spatial imagery, and the charts are in a graphical language. Furthermore, the internal representations may be influenced by a history of interaction with particular types of external representations. This is why chart depictions that are actually quite arbitrary seem natural to experienced navigators. Evaluating a flight path clearance requires a translation from one language of description to another.

Create Machine-Readable Description of Flight Path

Once the clearance has been received and at least minimally processed, we may ask what becomes of it in the cockpit system. Usually the next stop on the information trajectory is the altitude preselect window of the altitude alerter. This is the first step in creating a machine-readable description of flight path in the altitude-change task. Since altitude alerters are required by law in all civil turbojet aircraft, this portion of creating a machine-readable description is performed in traditional cockpits as well as in glass cockpits. Simply setting the altitude in the window of the altitude alerter provides another opportunity for redundant checking. Figure 7 shows the possible pathways of information in this simple operation. The checking that can be done by the PNF while setting the altitude alert is part of the task of doing the setting. That is, this is a closed-loop action in which the PNF adjusts the displayed value until no discrepancy remains between the selected altitude and the PNF's memory of the cleared altitude. Since the altitude alerter is located where it can be seen by both pilots, it is also possible for the PF to notice a discrepancy between the displayed value and the remembered value.

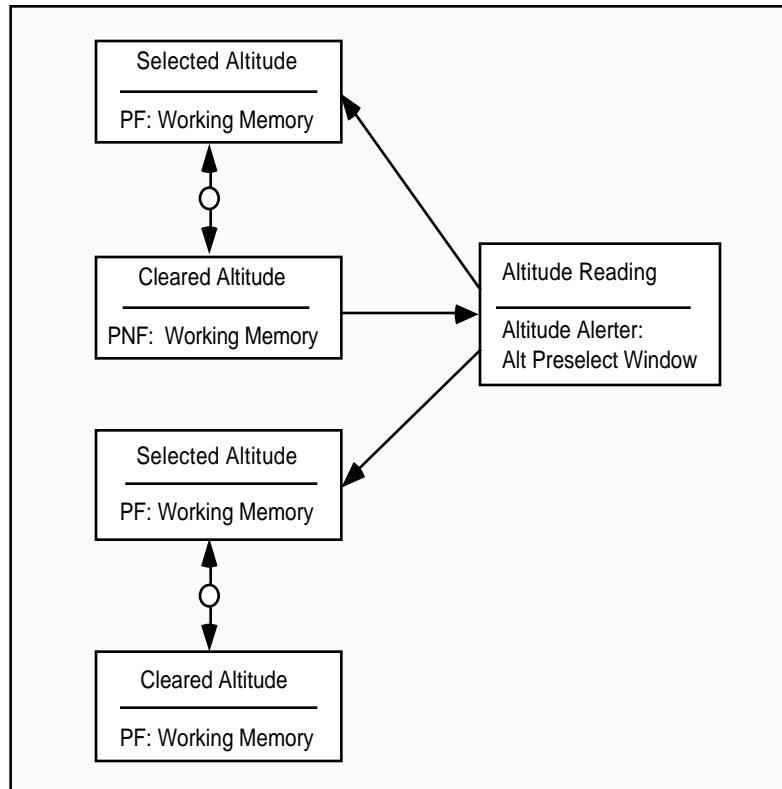


Figure 7. Information flow and checks during setting of altitude alert.

The representation of altitude in the altitude alerter often becomes the primary memory for the target altitude in the system. Pilots frequently seem to give this representation precedence over their own internal memories. Halverson (U. of California, San Diego, California; manuscript in work) reports an incident in which a pilot leveled off at the altitude shown in the altitude-alerter window even though he heard and correctly read back a clearance to a different altitude only moments earlier. In narrative No. 4, in a rather subtle example, a captain failed to question an erroneous setting of the altitude window. It seemed odd to him because it did not agree with what he thought he had heard and because the airplane's performance seemed unusual:

As we progressed, I noticed I was falling behind the required descent profile. Checking the altitude window, I saw 11,000 ft. I asked first officer if we had been cleared to 11,000 ft, thus cancelling the 13,000 ft restriction. He replied that the clearance was to cross the fix at 11,000 ft.

Yet the captain did not question it because the 11,000 ft altitude met his long-term expectations about the usual altitude of the airplane at this point in the flight:

Due to the turbulence and the fact that we normally are at 11,000 ft at that point in the profile, I failed to have the first officer verify our clearance before descending below 13,000 ft.

This led to an altitude bust.

Recall that one of the functions of the read-back is to permit the controller to be certain that the pilots intend to do what the controller has asked them to do. Even if the pilots have read back the clearance correctly, there is no guarantee that they have understood it correctly, or that they have taken the appropriate actions with respect to it. In particular, it is possible to hear and read-back a clearance and then either forget to reset or misset the altitude in the altitude alerter. This has been observed in our sample of reports and in a full-mission simulation (Halverson, C., U. of California,

San Diego, California; manuscript in work). One possible technological solution for this situation would be to downlink the altitude set in the alerter via data-link (Billings 1991). Like the read-back, this would permit the controller (or the ATC computer) to be certain that the autoflight system has been set to have the same intention as the pilots and that the pilots have the same intention as the controller. Downlinking the setting of the altitude alerter would provide another opportunity for error detection as shown in Figure 8. Incident No. 4 probably would have been detected earlier by the controller if the altitude preselect window setting had been downlinked to the ATC computer and to the controllers' display.

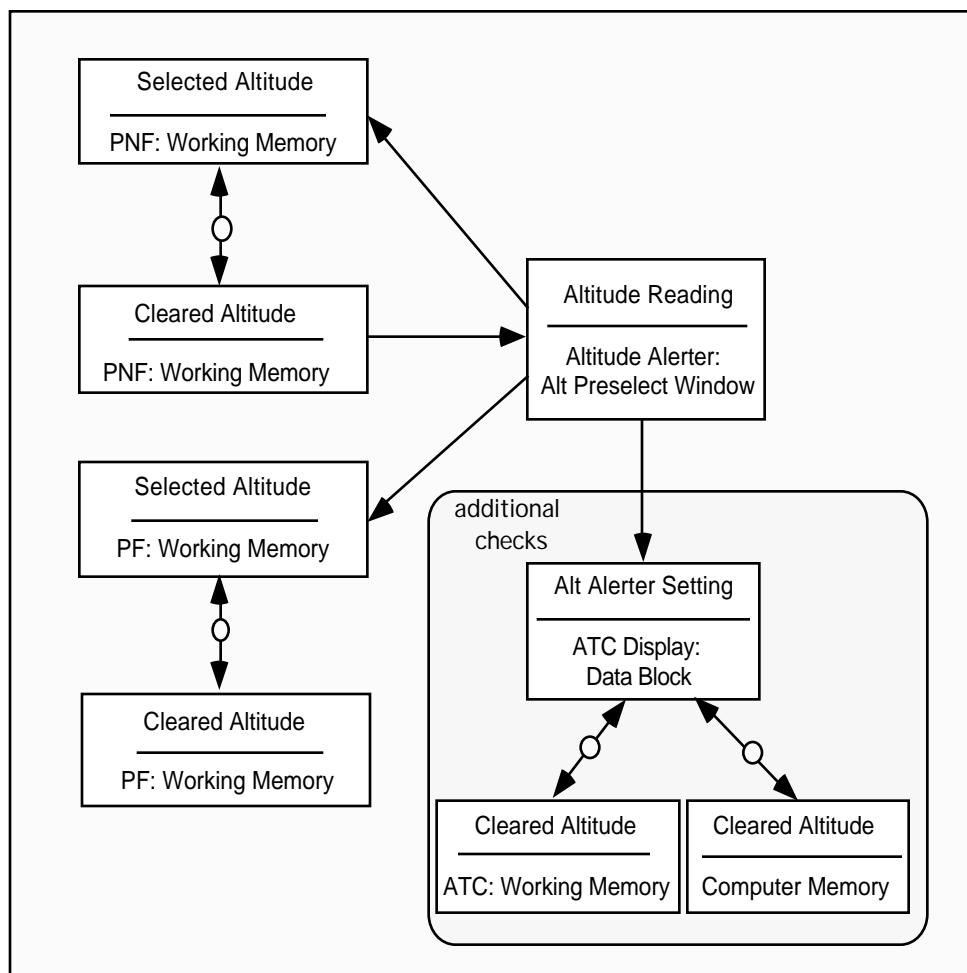


Figure 8. Additional check possible with downlinked value of the altitude alerter.

More Complex Machine-Readable Descriptions of Flight-Path

Flying always involves a relationship between a pilot's mental description of the flight path and the behavior of the airplane. With even the simplest autopilot, goals must be specified for the automation. Even if they are no more than "maintain current altitude and speed," these goals must be represented in a way that the machinery can make use of them. In more complex systems, such as those found in the glass cockpit, a much wider range of goals, subgoals, and constraints can be specified for the autoflight systems. In fact, aircraft with FMSs can now use a description of an entire flight to guide an airplane from initial climbout on one coast to touchdown and rollout on the runway on the opposite coast. If the crew of an automated airplane wishes the autoflight systems to fly the flight path specified in a clearance, the clearance information must be translated into a form

that the machine can utilize. The machine-readable description of the flight path must then be linked to the autoflight computers, and finally the autoflight systems must be linked to the control surfaces of the airplane.

We have already discussed the processes by which the information in the spoken clearance makes its way via the minds of the pilots to the altitude-alerter display. As we saw, there are many opportunities for error detection in the process of getting the number into the window. The display becomes a memory for the cleared altitude, but it may be more than that. In a glass cockpit, the altitude in the window may also be an element of a machine-readable description of flight path. The construction of such descriptions is the next major step in the process.

Consider narrative No. 2:

Center cleared us to make a heading change to intercept a radial and cross 30 DME prior to the VOR at 11,000 ft. I verified that we could make the restriction and observed the first officer program the computer. The new first officer made several errors in the fairly lengthy procedure which required a new waypoint, a radial to the waypoint, crossing restrictions at the waypoint as well as a turn to intercept the new course. Time became short, and as I attempted to assist him in correcting his errors, I inadvertently inserted the wrong waypoint. (I inserted the radial/altitude rather than the radial/DME as the fix.) I then assisted the first officer in accomplishing the procedure. I did not realize the error until we were leveling at the waypoint and approach control asked us if center had given us the crossing restriction.

This error emphasizes that the complex cockpit can, at times, be a hindrance rather than an advantage. A constant verification of data and inputs is required.

The description of the flight path that must be entered into the control and display unit (CDU) of the FMS is quite complicated. For example, the procedure described in this narrative would require the following 14 steps in a typical glass cockpit:

1. Push the heading select button on the MCP and slew the heading bug to the intercept heading. This disengages LNAV.
2. Enter 11000 in the altitude alerter window on the MCP.
3. Enter the VOR identifier on the CDU scratch pad.
4. Line-select the VOR identifier to line 1 left of the RTE LEGS page.
5. Enter the inbound course to the scratch pad (this will be the reciprocal of the VOR radial specified by ATC and the conversion will have to be made by the crew).
6. Line-select the specified inbound course to line 6 right of the RTE LEGS page.
7. Execute the modifications to this point by pressing the CDU EXEC button.
8. Line-select the VOR back to the scratch pad and add the characters “/-30” to create a waypoint 30 n. mi. prior to the VOR.
9. Line-select this new waypoint description on top of the VOR identifier in line 1 left.
10. Enter “/110” or “11000” in the scratch pad.
11. Line-select this to line 1 right (adjacent to the newly created waypoint).
12. If necessary, close the route discontinuity following the VOR waypoint by selecting a next on-route waypoint.
13. Execute the route modifications by pressing the CDU EXEC button.
14. Re-arm LNAV by pressing the LNAV button on the MCP.

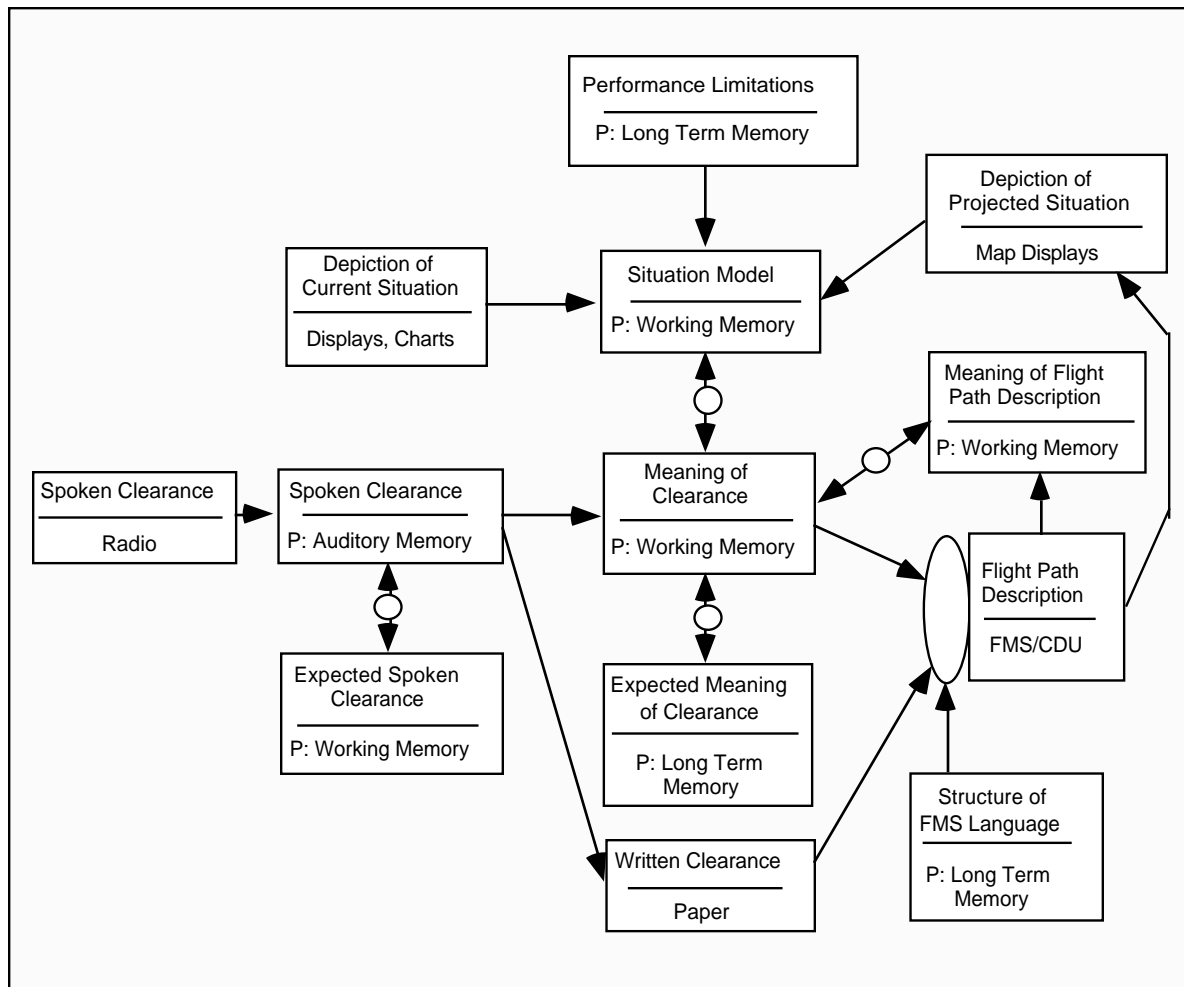


Figure 9. Information flow in programming the FMS.

The description that is entered into the FMC may be constructed from memory for the spoken clearance, from the memory for the meaning of the clearance, and/or from a written record of the clearance if one was made or from a combination of these (Fig. 9). The translation from the spoken, written, and spatial representations into the input language of the FMS is mediated by knowledge of the FMS language conventions. The entry of the description is executed through the MCP and the FMS/CDU. Every step in the procedure described above entails a cycle of activity in which the gulfs of execution and evaluation are bridged. The evaluation of character strings that are entered with the CDU keypad is made on the CDU itself. The evaluation of the horizontal component of the flight path description can be made graphically on the map display. No graphical display is provided for evaluating the vertical component of the flight path, which must be evaluated on the text-based CDU. This example shows that errors in the programming can be detected by the other pilot in the cockpit, who may be monitoring the programming activity, or by ATC when the airplane flies an unexpected path.

Our sample of altitude deviations contained several examples in which pilots appeared to have difficulty bridging the gulf of execution and translating from a straightforward clearance in ATC language to the flight path description required by the autoflight and flight-management systems. There were also examples of problems bridging the gulf of evaluation where an error was made in describing the clearance in the language of the FMS, but the crew was not aware that they had made an error in entering the clearance until queried by the ATC controller. The map display is

useful for detecting gross errors such as route discontinuities in the horizontal flight path but is not always helpful in detecting more subtle horizontal errors or errors in the vertical path. The current FMSs do not help the pilots check to ensure that what they have entered in the CDU matches the clearance that they received from ATC. Perhaps the FMS should translate what the pilot has input back into ATC language. This might help the pilot check that the description in the FMS matches the clearance from ATC.

Another gulf-of-execution problem was with the MCP for the autopilot. The FMS is powerful enough to allow almost any clearance to be entered given enough time, but some of the altitude deviations in our sample could probably have been avoided if the autoflight MCP allowed the pilot to input constraints on the aircraft's airspeed during a low-altitude climb and level-off.

Applying the concepts of the gulf of execution and gulf of evaluation to the problems of translating an ATC clearance into a flight path description in the FMS is another way of describing the pilot's complaint in Wiener's field study about the increased workload required to use the automatic equipment in busy terminal areas. The task of FMS programming takes time, is mentally absorbing, and requires pilots to go "heads-down." Also, if one is interrupted in the middle of a programming task, it is often difficult to determine where one was and how to resume the task. The existence of several different flight-path descriptions (i.e., languages) can make it difficult to describe the desired flight path and increases the likelihood that errors will be made in describing a flight path to the autoflight system.¹⁷

Select Mode of Operation

Once elements of the flight-path description have been entered in machine-readable form, there remains a large number of possibilities for using the autoflight systems to realize the goals expressed in the flight path description. Pilots may choose a combination of thrust, roll, and pitch modes for the airplane, or may choose to disregard autoflight system computations and fly on "raw data." Some modes may be armed, awaiting the existence of certain conditions before automatically engaging, while others are engaged. This range of alternatives is too complicated to permit a complete analysis here. Clearly, though, there are problems with the mapping from intention to action in this element of the task.

The selections available on the MCP form a language of mode specification. Every possible combination of settings on the MCP is a potential sentence in this language of mode specification. The problem may be stated generally as one of knowing how to say what is wanted in this abstract language. The gulf of execution must be bridged by determining which MCP settings will accomplish the mode-selection goals. There are also problems on the side of the gulf of evaluation. Instructors commonly complain that crews in training use the state of the MCP as part of the evaluation of the mode status, when this is not appropriate. It is appropriate to refer to the MCP to determine that the setting produced is the setting intended. In order to determine that the setting produced on the MCP has actually engaged or armed the modes intended, it is necessary to refer to the mode annunciators, which are not located on the MCP. These mode annunciators represent autoflight status in yet another language which introduces even more cognitive complexity to the cockpit.

Couple the Autoflight Computers to the Flight Controls

If an automated mode has been selected, the outputs of the autoflight computers can be coupled to the flight controls in two ways. First, outputs of the autoflight computers can be used to drive the flight director and the pilot can then use the manual flight-controls to produce the attitudes indicated by the flight director command bars. Second, the outputs of the autoflight computers can be sent to the autopilot servos that drive the flight control actuators. This second option is activated by engaging one or more autopilots. Engaging an autopilot is another pilot action that requires

¹⁷ The use of the term "programming" to describe this task is telling. Computer "programming" is notorious for being mentally absorbing and costly to interrupt.

modify any of the vertical modes on the MCP and there was no immediate change in aircraft behavior in the vertical dimension. The crew had to remember that the vertical path would be affected by disengaging the autopilot to stop the horizontal turn.

This kind of ASRS incident together with observations of pilot behavior in full-mission simulators has led us to hypothesize that the pilot's primary sources of information about the mode of the autoflight system are (1) the explicit selections made on the MCP and (2) the change in behavior of the aircraft. Pilots are taught that the only reliable indicator of the mode of the autoflight system is the mode annunciator display. However, the hypothesis that the mode annunciator display is used as only a secondary source of mode information provides a possible explanation for several mode-awareness problems that we have observed in ASRS reports and in full-mission simulations. People naturally learn expectations for a system's immediate response to their control inputs. Problems occur when control inputs cause only an immediate change in the mode annunciator display and cause no immediate change in the flight path of the aircraft.

In discussing design characteristics that make human errors more likely, Norman (1986) uses the term "side effect" to refer to events that occur but are secondary to the primary purpose of using the system. Side effects often result in errors. Avoiding problems caused by side effects requires the user to remember that the side effect has occurred or will occur and to remember to take action to avoid the usually unwanted consequences of the side effect. Often systems provide little support in helping users remember to check for side effects. For the flight crew in narrative No. 1, disabling the altitude capture was a side effect of turning off the autopilot.

DISCUSSION

In the preceding sections, narratives from ASRS reports were used to construct a descriptive framework of the cockpit as a *distributed cognitive* system. We have traced some of the paths information can take as it moves from the controller to the flight crew and to the autoflight system during the altitude-change task. We use the term *distributed* to emphasize that the cognitive processing that is important to performance at the system level is distributed across the various people in the system and the various artifacts and machines on the ground and in the cockpit. We use the term *cognitive* to emphasize that the system achieves its goals and controls and monitors the altitude-change task largely by processing information. The processing involves transmitting, storing, accessing, checking, and cross-checking information represented in various media.

We refer to the active human and machine components of this distributed information processing system as *agents*. In performing *tasks*, the agents in the system access information stored on various *media*, and perform tests and transformations on the information and output information and on actions to other media. The media in the system have different properties. Each medium has a limited number of ways to *represent* information. Media also differ in their *storage* properties. Information in an audio medium is only briefly available, whereas information represented on the altitude-alerter display is available until the setting is changed. Media also differ in their *accessibility* by other agents. The altitude-alerter display is centrally located and easily accessible by both pilots, the CDU of one pilot is accessible with difficulty, and information on a CDU page that is not displayed requires pilot input to the CDU to access. The pilots' memories are also modeled as media. These media have representations that are likely to be unreliable and are directly accessible by only one agent.

The key features of this system that make it error tolerant are that multiple agents have access to clearance information, and agents can redundantly check to ensure that tasks are performed and that information is processed correctly by the other agents in the system. Agents have expectations about how information represented in one medium (the spoken clearance) will appear in another medium (the MCP) and can use these expectations to detect and correct information processing errors.

The key features of this system that make it susceptible to breakdown are that it is constructed of fallible agents (people) that may occasionally make errors in performing tasks or processing

information.¹⁸ Because of distractions and lapses people may also occasionally fail to perform the checking and cross-checking tasks that are necessary to detect these errors and that give the overall system its error-tolerant properties. The terms “could,” “should,” and “may” were frequently used in the preceding section when describing checking and cross-checking tasks. Often these tasks are good practice, but do not have to be done in order that the altitude-change task be accomplished correctly.

Feedback and Error Tolerance

Fundamental to reliably performing any task without error by humans or by machines is the concept of feedback and checking. In general, feedback involves comparing the actual state of a controlled system with the desired state of the system. When a mismatch is detected, action is taken to reduce the mismatch. The need for feedback in aviation tasks is most apparent in continuous-control tasks such as continually detecting and compensating for path errors while manually flying an instrument approach. In continuous-control tasks, errors are an expected part of doing the task and displays are designed to clearly display errors and to make error detection and error compensation straightforward.

In the diagram of a continuous-control system (Fig. 11), feedback from the actual state to the desired state forms the “loop” that is the basis for the phrase “in the loop” that is used to describe the desirable state of being actively involved in the process and of having good situation awareness. Discrete tasks such as setting the cleared altitude into the altitude alerter are usually done without error and the need for feedback is not as apparent, but feedback and checking are necessary for any discrete or continuous task to be done consistently without error.

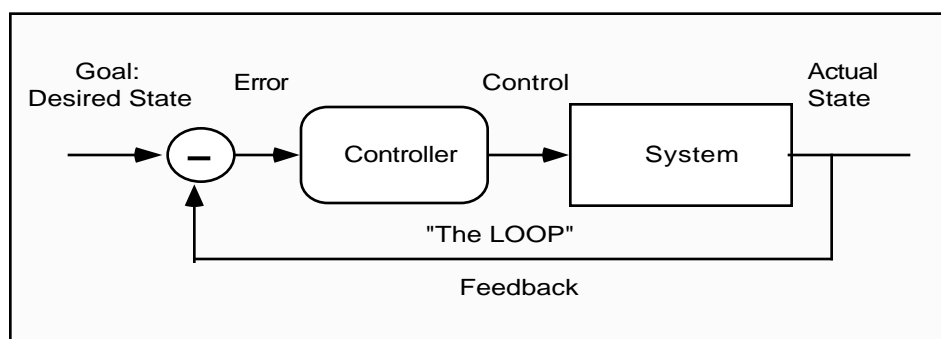


Figure 11. A simple servomechanism.

Distributed Access to Information

Multiple people and machines in the cockpit and on the ground are involved in performing the checking and cross-checking tasks associated with making an altitude change. Hutchins and Klausen (1991) describe how the pilots’ distributed access to information and each pilot’s expectations of what events should occur during an altitude change result in error detection and correction. Even though it is usually the PNF’s responsibility to communicate with ATC, the other crew members are usually also monitoring the ATC frequency in order to know the current clearance and to check for communication errors. Our sample of altitude-deviation incidents included several examples in which one pilot was off the ATC frequency and failed to pick up a communication error made by the other pilot. When an ATC clearance with a new altitude is received, one of the pilots enters it into the altitude-alerter window. This provides the command altitude for the autopilot, but just as importantly it provides information about the cleared altitude to

¹⁸ These errors may be caused by any number of factors including lack of attention, lack of knowledge, or equipment design that increases the likelihood of human errors. Equipment design that reduces the likelihood of human error is termed *error-resistant* design (Billings 1991).

the altitude alerter that allows machine monitoring of the altitude-change task and the same information to the other pilot to allow human monitoring of the autoflight system.

Distractions

Distractions can cause failures in both execution tasks and evaluative checking tasks. However, checking tasks seem to be much more likely to be disrupted or skipped because of a distraction (Monan 1979). There is rarely an obvious penalty for not performing a checking task. If the flight-path is described correctly and the automation functions correctly, the evaluative checking and cross-checking tasks do not have to be accomplished correctly, or at all, for error-free performance at the system level. The checking and cross-checking tasks provide the error-tolerant characteristics of the system that accomplishes the altitude-change task. In an error-tolerant system, checking tasks detect many errors and result in normal error free operation as viewed by an external observer. These tasks provide opportunities to detect and correct errors. They provide the redundancy necessary for the normal externally observable error-free execution of the altitude-change task.

Checking tasks that provide error-tolerant properties seem to be particularly susceptible to not being done or to being done incorrectly as a result of allocating attention to apparently more important tasks that can not be done correctly without focused pilot attention. However these interrupting “primary” tasks may actually be less important in terms of overall safe operation than the checking or cross-checking task. Pilots might even agree that performing the checking task is of higher priority than performing a primary task, but their behavior may not always accord with those priorities. For example, most pilots would agree that making the “thousand-feet-to-go” call is more important than talking to ATC or to their company, but we are sure that observations of cockpit behavior would show many instances when these “lower” priority tasks take precedence over more important safety-of-flight checking and cross-checking tasks.

Checking tasks are good practice, but may not necessarily be prescribed behavior. In a desire to reduce the number of altitude deviations in its fleet, one air carrier has formally implemented explicit crew procedures for making altitude changes (Sumwalt 1991). For example, the following is the mandated procedure for receiving a clearance and setting the altitude window:

Pilot accepting clearance will set altitude alerter/MCP (normally pilot not flying). S/he will announce and point to the new altitude until the other pilot (normally pilot flying) observes and repeats the assigned altitude. Cross cockpit verification is a must. LOOK AT IT.

A key feature of these procedures is that they mandate how and when checking and cross-checking tasks are to be accomplished. The procedures also require overt responses in tasks that otherwise could be done with only an eye scan. The overt verbal and gestural responses help insure that each pilot knows that the other pilot is aware of the new assigned altitude. The fact that each pilot knows that the other pilot knows the cleared altitude improves the likelihood of detecting errors. The explicit nature of the procedures should also improve each pilot’s expectations of how the other pilot should behave during the altitude-change task. Good expectations are necessary for detecting errors by observing deviations from expected behavior. If it is not clear what the other crew member is supposed to do, it is difficult to know if the other crew member’s behavior is correct or incorrect.

When check and cross-check tasks are not accomplished, tasks may still be accomplished correctly but they will be accomplished with less than the desired amount of redundancy. Crews may not be sensitive to the implications of the loss of these redundant checks on system safety.

Intelligent Feedback

A controversial question concerning cockpit design is whether the visual approaching-altitude alert should be accompanied by an aural signal. An early study on altitude deviations (Factors . . . 1977) noted that many pilots reported that the aural alert sounded so often that they frequently did not hear it. That report recommended that the aural alert be made optional in future aircraft. Now ASRS reports are received that blame reported altitude deviations on the lack of an aural alert.

Lyddane (1991) has proposed providing the aural approaching-altitude signal only if the altitude that is set in the altitude alerter will be violated if the pilot does not take immediate action to reduce the aircraft's vertical rate. The alert would be triggered based on a linear relationship between vertical rate and altitude to go. It would be essentially a ground-proximity-warning system¹⁹ (GPWS) for altitude deviations. An alert would not sound if the autoflight system performed the altitude change in the usual, expected way. Such an alert would be triggered on only the rare occasions that positive action was required to prevent a deviation. Unlike the current altitude-deviation warning, this alert would sound in time to prevent the deviation. This is the kind of intelligent feedback (Norman 1990) that systems must provide to support the flight crew in their more supervisory tasks. Many autoflight systems would not rate well on cockpit resource management (CRM) criteria for being an effective subordinate.

RECOMMENDATIONS

ASRS reports provide a window on problems that are occurring in today's aviation system. They are useful in identifying problems and stimulating ideas for correcting current problems.

Procedures

Procedures for doing the checking and cross-checking tasks during the altitude-change task should be formalized. A good example of how one airline has done this is described in the paper by Sumwalt (1991). Error tolerance depends on distributed access to information by "agents" that have expectations about correct behavior. Explicit procedures help by allowing crew members to have stronger expectations about the behavior of other crew members and by providing explicit behavioral signs so that each pilot knows that the other pilot knows the current cleared altitude.

Cockpit procedures should be designed to minimize the amount of time when only one pilot is on the radios. Unless both pilots hear the clearance, they cannot redundantly monitor the communication for errors.

Training

Several ASRS reports from pilots of glass-cockpit aircraft indicate that pilots could use more training on how to use the many features of the complex autoflight and FMSs in glass-cockpit aircraft. A specific recommendation is that pilots be provided with a selection of ASRS reports that illustrate the kinds of problems that can occur when using autoflight equipment.

Curriculum designers could use ASRS reports to determine the kinds of problems that pilots are having in line operations with autoflight and FMSs and then tailor initial and recurrent training to emphasize these problem areas.

¹⁹ The ground proximity warning system (GPWS) provides a predictive alert, based in part on radar altitude and vertical velocity, that inadvertent contact with the ground is imminent.

New Equipment: Datalink

The use of data-link for ATIS, company communications and weather information should reduce communication errors and increase the time during which all pilots are monitoring the ATC frequency.

The use of data-link has some potential advantage for the display of ATC clearances. Datalinked clearances could be displayed on a printed medium instead of through an auditory medium. Printed media have the advantage that information endures and, unlike with auditory media, pilots can control when they will attend to the information. The clearance information is also in a computer-readable format that could be input directly to the autoflight and flight management systems thus reducing entry time and the possibility of keying errors. However, the use of data-link for clearance delivery will require careful attention to the design of cockpit procedures to insure that when a visual medium replaces the current auditory media all crew members are aware of ATC clearance communications. Procedures such as that reported by Sumwalt may be even more important in a data-link environment than in today's largely voice communications system.

The use of data-link to close the intent loop between air and ground should be investigated. By downlinking the altitude that the pilot has set into the altitude alerter and displaying it on the controller's display, it is possible for the controller (or a ground computer system) to monitor redundantly that the flight crew has received and set the altitude alerter with the correct altitude. This is an example of the communication-of-intent information between air crew and air-traffic controller that is advocated by Billings (1991).

Future Design Possibilities

More intelligent feedback

The issue of whether to provide an aural alert when the aircraft is approaching the altitude set in the altitude alerter should be revisited. A more intelligent audio alert, such as that proposed by Lyddane (1991) that only sounds if an altitude deviation is imminent, would appear to alert pilots to an approaching altitude deviation without either sounding so often in routine situations that pilots start to ignore it or come to rely on it as their primary source of altitude information.

Displays

The gulfs of execution and evaluation should be reduced with electronic map enhancements such as the green arc, DME circles around navigational fixes, and energy circles. These allow pilots to control and monitor vertical profiles with little or no FMS programming. Research should also be done on a vertical analog of the horizontal map display to help the pilot access, plan, and monitor for vertical path errors (Baty and Watkins 1979; Fadden et al. 1991).

Primary flight displays

Primary flight displays should be developed that make busting an altitude perceptually dramatic. With current displays almost nothing visual happens. The altitude-alert system provides a "dimly lit light" (narrative No. 7). Otherwise, the task of monitoring the level-off involves comparing the altitude displayed on the altimeter with the reference value stored in a different location. The altitude bugs on the tape altimeters are a step in the right direction. The tunnel-in-the-sky display would seem to provide a more explicit indication that the aircraft has busted or is about to bust its altitude (Grunwald 1981; 1984). Busting an altitude should be as perceptually dramatic as failing to make a turn planned in your automobile on an expressway.

CONTINUING RESEARCH

In several ASRS reports, pilots in our sample mentioned fatigue caused by long duty days and insufficient sleep as a factor contributing to reported altitude deviations. There is an on-going research program at NASA Ames on the effects of flying schedules on flight crew sleep loss and on pilot performance. Studies of ASRS incident reports are being performed that focus on the link

between pilot fatigue and aircraft incidents. A full-mission simulator study is planned to evaluate flight crew performance after a multiday mission.

A computational model of the information processing in a transport cockpit during the descent and approach phases of flight is being constructed. This model will focus on the flow of information among the various people and machines in the system. The agents in the model will have behavioral expectations that allow them to monitor the tasks and information processing of the other agents and machines in the system. This computational model will be used to analyze the information flow during altitude changes made in well-documented full-mission flight simulations.

CONCLUSION

In this report we took an information processing approach to describing the system that performs the altitude-change task. Our unit of analysis was the cockpit system, not the individual pilot. We were concerned with how the system processes, stores, and makes decisions as it processes clearance information. Narratives from altitude-deviation reports were used to help construct a descriptive model of the cockpit as an information processing system. Incident data are particularly useful in helping to reveal how the system processes information.

We focused on describing the error-tolerant properties of the system and why breakdowns occasionally occur. An error-tolerant system is a system that is capable of detecting and correcting its internal processing errors. The cockpit system consists of two or three pilots supported by autoflight, flight-management, and alerting systems. These human and machine agents have distributed access to clearance information and perform redundant processing of information. Errors can be detected as deviations from either expected behavior or as deviations from expected information. Breakdowns in this system are possible because the checking and cross-checking tasks that give the system its error-tolerant properties may not be performed owing to distractions, fatigue, or other task demands. Monitoring and checking tasks are not required in order to accomplish the altitude-change task, and there is almost never a performance penalty for not performing checking tasks.

Describing the cockpit as a single information-processing system is useful because the system behaviors of interest are not determined solely by the behavior of the humans in the system. Altitude deviations in transport aircraft are usually the result of multiple errors. The information processing analysis also highlights the variety of languages and media used to describe the flight path as clearance information moves through the cockpit system.

The properties of the error-tolerant systems provide a justification for several of our recommendations. More explicit procedures allow each pilot to have better expectations about what tasks the other pilot will perform and the information that the other pilot knows. More effective training should also provide pilots with better expectations of how the autoflight and FMSs work. In addition to reducing errors, this knowledge should allow pilots to be more effective in detecting programming errors. The expanded use of data-link for nonATC information would minimize the time that one pilot is unable to monitor the communications with the air-traffic controller. The use of data-link for clearances will result in a more persistent error-free medium, but the change to data-link for clearance delivery will require new cockpit procedures to insure that both pilots are aware of the current clearance. Using data-link to downlink selected values, such as the value in the altitude-alterer window, offers a channel for closing the intent loop between air and ground. This downlink would provide a redundant means for the controller (or a ground system) to monitor the altitude to which the crew is intending to fly. The recommendation for a GPWS-like altitude alert offers the potential of improving the altitude alterer's monitoring performance. With the predictive logic proposed by Lyddane, the system could provide an alert as soon as the flight path deviated from the expected flight path to capture the selected altitude.

Finally the information processing approach led us to the observation that there are several sources of autoflight mode information in the cockpit and then to the hypothesis that mode-control actions on the MCP and the aircraft's behavior may be more salient clues to autoflight mode than the flight mode annunciator display.

Over 8,000 altitude deviations are reported to the ASRS each year—nearly one an hour. It is hoped that analyzing the way information is processed by the cockpit system will lead to procedural or design changes that will help reduce the number of these incidents.

National Aeronautics and Space Administration
Ames Research Center
Moffett Field, California 94035-1000
199

APPENDIX: Glossary

Aviation terms

ACARS: ARINC communications and address reporting system: digital communications system used primarily for aircraft-to-airline messages

ARINC: Aeronautical Radio, Inc.

ASRS: aviation safety reporting system: aviation incident reporting system run by NASA for FAA

ATC: air traffic control

ATIS: automated terminal information service: recorded voice message that provides weather and airport services information

CDU: control display unit: pilots' interface to the FMS

CRM: cockpit resource management

DME: distance measuring equipment: ground navigational aid that can provide display of distance to selected ground navigational radio transmitter

EFIS: electronic flight instrument system

FAA: Federal Aviation Administration

FAR: Federal Aviation Regulations: federal rules under which flight operations are conducted

Fix: position in space usually on aircraft's flight plan

FL 310: for example, FL310 is an altitude 31,000 ft above sea level; used for altitudes above 18,000 ft

FMA: flight mode annunciator: display on or near the PFDs of the current modes of autoflight system

FMC: flight-management computer

FMS: flight-management system

FO: first officer

GPWS: ground proximity warning system: warns of inadequate separation from ground and excessive sink rate close to ground

IAS: indicated airspeed

ILS: instrument landing system: uses precision localizer and glide-slope radio transmitters near a runway to provide landing approach guidance

LNAV: lateral navigation: provides computer description of aircraft's planned lateral flight path that can be tracked by the autoflight system; lateral path can be shown on map display.

MCP: mode-control panel: pilots' interface to the autoflight system; usually located centrally just below cockpit glare shield

NASA: National Aeronautics and Space Administration

PF: pilot-flying

PFD: primary flight display

PNF: pilot-not-flying

PPOS: present position

RNAV: area navigation: generic acronym for any device capable of aircraft guidance between pilot-defined waypoints

VNAV: vertical navigation: provides computer description of aircraft's speed and altitude that can be tracked by autoflight system

VOR: very-high-frequency omnidirectional radio range: ground navigational aid that can provide display of aircraft position relative to course through selected ground navigational radio transmitter

Waypoint: position in space usually on aircraft's flight plan

Cognitive science terms

Action: goal directed behavior performed by an agent

Agent: any system component that executes tasks that process information

Attention allocation: process by which an agent samples information on media in the environment

Cognitive tasks: tasks that achieve their goals by processing information

Distraction: an event that causes attention to move away from task to which agent should continue to attend

Distributed cognition: term used to emphasize that cognitive tasks that are important to performance at system level are distributed across various people in the system and various artifacts and machines

Error-resistant system: system designed to minimize occurrence of errors

Error-tolerant system: system that can detect errors and correct them before they result in undesirable consequences

Expectation: representation of a state predicted to occur in a given situation; agents can compare expectations to what actually happens and thereby detect possible processing errors

Information: describes the relation between a structure (in the world) and some agent that interprets that structure; world is full of structure, but the only structures that are informative are those capable of affecting the behavior of the agent that is interpreting those structures

Information processing: tasks that depending on the situation, transform information represented in one medium to information represented in a different medium

Media: any physical things that can be used to represent information

Mediated information: information that has been transformed from its primary source

Mistake: an error in which the executed action is the intended action but the intended action is incorrect

Representation: structure in the world used to express information on a medium

Slip: an error in which the executed action is not the intended action

Side Effect: change in a machine state that is a byproduct of an agent using the machine to accomplish some task

Task: actions accomplished by agents to satisfy goals

REFERENCES

- Baron, S.; Kruser, D. S.; and Huey, B. M.: Quantitative Modeling of Human Performance in Complex Dynamic Systems. National Academy Press, Washington, D. C., 1990
- Baty, D. L.; and Watkins, M. I.: An Advanced Cockpit Instrumentation System: The Coordinated Cockpit Display. NASA TM-78559, 1979.
- Billings, C. E.: Human-Centered Aircraft Automation: A Concept and Guidelines. NASA TM-103885, 1991.
- Billings, C. E.; and Cheaney, E. S.: The Information Transfer Problem: Summary and Comments. NASA TP-8567, 1981.
- Curry, R. E.: The Introduction of New Cockpit Technology: A Human Factors Study. NASA TM-86659, 1985.
- Degani, A.; Chappell, S. L.; and Hayes, M. S.: Who or What Saved the Day: A Comparison of Traditional and Glass Cockpits. Sixth International Symposium on Aviation Psychology, Columbus, Ohio, 1991.
- Factors Associated with Altitude Overshoots, Excursions, and Undershoots. NASA TM-X-3546, 1977.
- Fadden, D. M.; Braune, R.; and Wiedemann, J.: Spatial Displays as a Means to Increase Pilot Situation Awareness. Pictorial Communication in Virtual and Real Environments, S. R. Ellis, M. K. Kaiser, and A. J. Grunwald, eds., Taylor and Francis, New York, 1991, pp. 172-181.
- Grunwald, A. J.: Experimental Evaluation of a Perspective Tunnel Display for Three-Dimensional Helicopter Approaches. J. Guidance Contr., vol. 4, Nov.-Dec. 1981, pp. 623-631.
- Grunwald, A. J.: Tunnel Display for Four-Dimensional Fixed-Wing Approaches. J. Guidance Contr. Dyn., May-June 1984, pp. 369-377.
- Human Factors Associated with Altitude Alert Systems. NASA TM-78511, 1978.
- Hutchins, E. L.; Hollan, J. D.; and Norman, D. A.: Direct Manipulation Interfaces. New Perspective on Human-Computer Interaction, D. A. Norman and S. Draper, eds., Lawrence Erlbaum Associates, Hillsdale, N. J., 1986, pp. 87-124.
- Hutchins, E. L.; and Klausen, T.: Distributed Cognition in an Airline Cockpit. Cognition and Communication at Work, D. Middleton and Y. Engestrom, eds., Cambridge University, 1992.
- Lyddane, G. H.: Aircraft Design and Certification Issues. Challenges in Aviation Human Factors: The National Plan, 1991, pp. 93-95.
- Monan, W. P.: Distraction: A Human Factor in Air Carrier Hazard Events. NASA TM-78608, 1979.
- Norman, D. A.: Categorization of Action Slips. Psych. Rev., vol. 88, 1981, pp. 1-15.
- Norman, D. A. ed.: Cognitive Engineering. Lawrence Erlbaum Associates, Hillsdale, N. J., 1986.
- Norman, D. A.: The "Problem" of Automation: Inappropriate Feedback and Interaction, Not "Over-Automation." Phil. Trans. R. Soc. London, vol. 327, 1990, pp.585-593.
- Reason, J.: Human Error. Cambridge University Press, 1990.
- Reason, J.; and Mycielska, K.: Absent Minded: The Psychology of Mental Lapses and Everyday Errors. Prentice-Hall, Englewood Cliffs, N. J. 1982.
- Sumwalt, R. L.: Eliminating Pilot-Caused Altitude Deviations: A Human Factors Approach. Sixth International Symposium on Aviation Psychology, Columbus, Ohio, 1991.
- Wiener, E. L.: Human Factors of Cockpit Automation: A Field Study of Flight Crew Transition. NASA CR-127333, 1985a.
- Wiener, E. L.: Cockpit Automation: In Need of a Philosophy. Society of Automotive Engineers, TP 851956, 1985b.

Wiener, E. L.: Fallible Humans and Vulnerable Systems: Lessons Learned from Aviation.
Information Systems: Failure Analysis, N. Asi, ed., Springer-Verlag, Berlin, 1987.
Wiener, E. L.: The Human Factors of Advanced Technology (Glass Cockpit) Transport Aircraft.
NASA CR-177528, 1989.